

Learning, verifying, and erasing errors on a chaotic and highly entangled programmable quantum simulator

Thesis by
Adam Lawrence Shaw

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended April 11, 2024

© 2024

Adam Lawrence Shaw

All rights reserved

Acknowledgements

Thanks go first and foremost to my advisor Manuel Endres. Manuel took me on despite my having what felt like no experimental, AMO, or quantum experience – in our first meeting I asked what type of dil fridge or cryo system was used to keep the atoms ‘ultracold’ – and every day since he has been a model to me for the precision, rigor, and creativity all scientists should aspire for. With his patented *step-by-step*TM approach to science, Manuel has a remarkable eye for plotting a course through the stormy seas of research, and for (usually) avoiding the murky waters of *verschlimmbessern*. I am deeply grateful to him for guiding me to become a better physicist, and for mentoring me throughout the successes of the last six years.

I would be remiss if I didn’t highlight Soonwon Choi, my theoretical pseudo-advisor from the Caltech of the East. Soonwon began collaborating with our group in early 2020, an alliance which has been incredibly fruitful and continues to this day. Further, I owe a huge amount of the success of this thesis to Soonwon’s students with whom I’ve shared tight partnerships. Daniel Mark has the unerring ability to be able to provide answers to practically any question I pose him, and to be able to invent theoretical proofs and numerical analysis techniques seemingly from thin air. His deep analytical knowledge hugely elevated many projects. Zhuo Chen was my constant companion on our approximate benchmarking project, which was the most complex of my PhD by orders of magnitude; he literally enabled the entire work through his peerless numerical abilities, his dedication, and his patience when I continually pestered him for more (and bigger) simulations.

In many ways, my Ph.D. was (in)conveniently trifurcated by the Covid-19 pandemic, with three somewhat distinct eras.

In the first – from when I joined the lab in 2018 until early 2020 – I was graced with incredible mentors in the lab. Alex Cooper-Roy and Jake Covey were my first postdocs: both provided constant mentorship and dramatically accelerated my development in the lab.

A special thank you is due to Ivaylo Madjarov, who was my experimental co-pilot, my sole graduate student compatriot, and an altogether stellar example of a physicist. When the pandemic hit, Ivaylo deferred to let me be the sole experimentalist to return to the lab; I thank him for trusting the experiment to me, though it was made easy by the fact that he got distracted by online chess. I want to reiterate and emphasize my appreciation for all those who came before me (including those I didn't overlap with) for building a firm foundation for the experiment which I've relied on. I often tell tours of our lab that our experiment is built so stably that I've only touched the MOT mirrors once in the last six years, for which I am very grateful.

Then came the pandemic. During this time I returned to lab as the sole experimentalist for a little over half a year. However, despite my solitary time in lab, I was certainly not without support. Besides Ivaylo's continued and steadfast guidance, I was later joined by Xin Xie who brought a huge wealth of AMO acumen, and who led upgrades of core components of the experiment. Further, for the entirety of the pandemic (and beyond) I had with me one of the finest partners imaginable in Joonhee Choi, who contributed deep theoretical knowledge, bountiful numerical skills, and a compassion to answer any question one might ask. It is little surprise given our long and ongoing collaboration that I have exchanged more Slack messages with him than with all other members of the lab, combined, times ten.

Now at the close of my PhD, we are in the third and final era, when I transitioned from being the sole experimentalist to handing off the reins to the next generation. Thanks go first to postdocs Ran Finkelstein and Pascal Scholl, who joined me directly on the experiment, collaborated closely with me on multiple projects, and generally expanded our scientific horizons. Then, there is the cadre of graduate students who have entered the lab since I joined, and who I know will continue to elevate it beyond my capabilities as I now depart. First, there is Hannah Manetsch (yes, from the celebrated Manetsch et al.) who almost always said yes to getting ice cream with me, and has been my constant (business) friend. To Gyohei Nomura (the star athlete) and Elie Bataille (the star student), I say 'bonjour' and '山の刀はどこですか' for putting up with my many linguistic butchering brilliancies. Then, to Richard Tsai and Xiangkai Sun, my most worthy successors, I wish you the best of luck managing lab goblins, as they say; though my time co-running the experiment with you was limited, I am assured in your abilities, and am excited to see where you take the experiment from here.

Besides all those mentioned explicitly, I am incredibly grateful to have been able to have collaborated with, and learned from, a broad swathe of people throughout my PhD, including but not limited to: Hannes Pichler, Andreas Elben, Kon Leung, Xudong Lv, Franz Pfanner,

Zunqi Li, Dariel Mok, Vlad Schkolnik, Jordan Cotler, Su Direkci, Arian Jadbabaie, Robert Huang, Federica Surace, Wenjie Gong, Fernando Brandão, Tai Hyun Yoon, Andrew Ivanov, Jasen Zion, Anant Kale, and many more. All have brought unique perspectives and crucial contributions to a wide variety of research fronts. Special thanks go to the AMO community at Caltech for being generous with both their advice and their optics.

Outside of Caltech, I've been extremely lucky in the past six years to have a wealth of familial support during my time in graduate school. To my brother, aunts, uncles, cousins, and grandparents – both around Los Angeles and beyond – thank you for your interest in my work, for taking me out to dinner, and for letting me take home the leftovers. In particular, thank you to my godfather Corwin Booth: he inspired my love of physics, and while I have strayed from the ideal of his actinides, he has continually supported me all the same. Deepest appreciation to Jane Wu for being my constant partner in the never ending struggle to keep sane during the ordeal that is grad school.

Finally, thank you to my parents. To my father Steven Shaw, who was one of the most creative people I knew, who fostered my love of learning, and who taught me to never lose sight of the things which are most important. While he passed away before seeing me take up the mantle of 'Dr. Shaw,' I am assured in his love, always. To my mother Heida Shaw, who taught me about atoms (and who had the foresight to name me Adam) and who supports me unconditionally in everything I do: I hope that I can reciprocate a fraction of the love she has always shown me. That's relatively close to the same acknowledgement as I put for her in my undergraduate thesis, which hopefully reinforces how unconditional and immutable her support has remained. This thesis is dedicated to both of them, for bringing me this far, and for always answering why.

In deference to my kindergarten self, I reiterate that I am thankful for the solar system.

Abstract

Controlled quantum systems have the potential to make major advancements in tasks ranging from computing to metrology. In recent years, quantum devices have experienced tremendous progress, reaching meaningful, intermediate-scale sizes and demonstrating advantage over their classical counterparts. Still, sensing, learning, verifying, and hopefully mitigating errors in these systems is an outstanding and ubiquitous challenge facing all modern quantum platforms.

Here we review and expound upon one such platform: arrays of Rydberg atoms trapped in optical tweezers. We demonstrate several key advancements, including the first experimental realization of erasure conversion to prepare two-qubit Bell states with a fidelity in excess of 0.999, and to cool atoms to their motional ground state. We further showcase the tools of universal quantum processing via arbitrary single-qubit gates, fixed two-qubit gates, and mid-circuit measurement, and discuss applications of these techniques for metrology and computing.

Then, we turn to the many-body regime, generating highly entangled states with up to 60 atoms through analog quench dynamics. We reveal the emergence of random behavior from unitary quantum evolution, and uncover a universal form of quantum ergodicity linking quantum and statistical mechanics. We exploit these discoveries to verify the global many-body fidelity and then realize practical applications like parameter estimation and noise learning. Finally, we compare against both state-of-the-art quantum and classical processors: we introduce a new proxy for the experimental mixed state entanglement which is comparable amongst all quantum platforms, and that reflects the classical complexity of quantum simulation.

Published content and contributions

1. **A.L. Shaw[†] (co-first author)**, D.K. Mark[†], J. Choi, R. Finkelstein, P. Scholl, S. Choi, M. Endres. Universal fluctuations and noise learning from Hilbert-space ergodicity. arXiv:2403.11971 (2024).
A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.
2. R. Finkelstein[†], R.B.S. Tsai[†], X. Sun, P. Scholl, S. Direkci, T. Gefen, J. Choi, **A.L. Shaw**, M. Endres. Universal quantum operations and ancilla-based readout for tweezer clocks. arXiv:2402.16220 (2024).
A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.
3. P. Scholl[†], **A.L. Shaw[†] (co-first author)**, R. Finkelstein[†], R.B.S. Tsai, J. Choi, M. Endres. Erasure-cooling, control, and hyper-entanglement of motion in atom arrays. In revision at Science, arXiv:2311.15580, 2023.
A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.
4. **A.L. Shaw[†] (co-first author, corresponding author)**, Z. Chen[†], J. Choi[†], D.K. Mark[†], P. Scholl, R. Finkelstein, A. Elben, S. Choi, M. Endres. Benchmarking highly entangled states on a 60-atom analog quantum simulator. Nature 628, 71-77 (2024). doi:10.1038/s41586-024-07173-x
A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.
5. **A.L. Shaw[†] (co-first author)**, R. Finkelstein[†], R.B.S. Tsai, P. Scholl, T.H. Yoon, J. Choi, M. Endres. Multi-ensemble metrology by programming local rotations with atom movements. Nature Physics 20, 195–201 (2024). doi:10.1038/s41567-023-02323-w
A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.
6. P. Scholl[†], **A.L. Shaw[†] (co-first author)**, R.B.S. Tsai, R. Finkelstein, J. Choi, M. Endres. Erasure conversion in a high-fidelity Rydberg quantum simulator. Nature

622, 273-278 (2023). doi:10.1038/s41586-023-06516-4

A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.

7. **A.L. Shaw**, P. Scholl, R. Finkelstein, I.S. Madjarov, B. Grinkemeyer, M. Endres. Dark-state enhanced loading of an optical tweezer array. *Physical Review Letters* 130, 193402 (2023). doi:10.1103/PhysRevLett.130.193402

A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.

8. J. Choi[†], **A.L. Shaw[†] (co-first author)**, I.S. Madjarov, X. Xie, R. Finkelstein, J.P. Covey, J.S. Cotler, D.K. Mark, H.Y. Huang, A. Kale, H. Pichler, F.G.S.L. Brandão, S. Choi, M. Endres. Preparing random states and benchmarking with many-body quantum chaos. *Nature* 613, 468-473 (2023). doi:10.1038/s41586-022-05442-1

A.L. Shaw contributed to the experimental, numerical, and theoretical advancements.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	vi
Published content and contributions	vii
Chapter 1: Introduction	1
1.1 Background	1
1.2 The challenge of quantum science	3
1.3 Quantum platforms	4
1.4 Outline	5
I Rydberg atom arrays	7
Prelude	8
Chapter 2: Experimental setup and error modeling	9
2.1 Trapping and controlling single atoms	9
2.1.1 The Rydberg Hamiltonian	11
2.2 Strontium	11
2.3 Errors for Rydberg atom arrays	15
2.3.1 Summary	16
2.3.2 Laser errors	19
2.3.3 Temperature errors	26
2.3.4 State preparation and measurement (SPAM) errors	29
2.3.5 Spontaneous decay errors	31
Chapter 3: Erasure conversion	35
3.1 Locating errors	36
3.2 An atomic erasure conversion implementation	37
3.3 Case study: Improving Bell state generation	39
3.3.1 Protocol	40
3.3.2 Results	41
3.3.3 Error modelling	43
3.3.4 Extraction of the Bell state fidelity	45
3.3.5 Interaction limitation for Bell fidelity	49
3.4 Case study: GHZ state preparation	50
3.4.1 Overview	50
3.4.2 Learning from erasure errors	54
3.4.3 Statistics reduction due to erasure-excision	56
3.5 Erasure correction cooling	57
3.5.1 Overview	58

3.5.2	The mechanism of erasure correction cooling	62
3.5.3	Cooling simulations	63
3.5.4	Fundamental limits simulations	65
3.5.5	Erasure correction cooling on other platforms	67
3.5.6	Erasure cooling in the axial direction	68
3.6	Dark-state enhanced loading	69
3.6.1	Overview	69
3.6.2	Method	71
3.6.3	Results	73
3.6.4	Discussion	76
Chapter 4:	Tools of a universal quantum processor	78
4.1	Local addressing with atom shifts	79
4.1.1	Light-matter interaction Hamiltonian	84
4.1.2	Operation fidelities	84
4.2	Mid-circuit readout	86
4.2.1	Preparing motional superposition states	87
4.2.2	Mid-circuit readout with motional state shelving	89
4.2.3	Motional coherence and motional-echo sequence	90
4.3	Entanglement and hyper-entanglement	91
4.3.1	Entanglement	92
4.3.2	Hyper-entanglement	94
4.3.3	Details of the two-qubit Rydberg gate	95
4.3.4	Parity oscillations for the motional Bell states	95
4.4	Application: Metrology	96
4.4.1	Quadrature readout	97
4.4.2	Multi-ensemble metrology	100
4.4.3	Estimating laser phase prior width	104
4.4.4	Fitting the phase-slip probability	104
4.4.5	Limits in multi-ensemble metrology	106
Coda		108

II Randomness and verification 110

Prelude	111
Chapter 5: Uncovering emergent randomness	113
5.1 Haar-random state ensembles	113
5.1.1 Observables of random ensembles	116
5.1.2 Quantum state k-designs	120
5.1.3 Why does randomness matter?	124
5.1.4 Relation to spherical designs	125
5.2 The projected ensemble	127
5.2.1 Moments of the projected ensemble	131
5.2.2 Comparison to quantum state k-designs	133
5.2.3 Finite temperature projected ensembles	135
Chapter 6: Many-body fidelity benchmarking	139

6.1	Estimating fidelity	140
6.1.1	The linear cross-entropy benchmark	141
6.1.2	Caveats of cross-entropy benchmarking	142
6.2	Building a better cross-entropy	145
6.2.1	Experimental results	147
6.2.2	Data analysis	149
6.2.3	Calculating the entanglement saturation time	150
6.2.4	Derivation of the fidelity estimator	150
6.2.5	Fidelity estimation in the Rydberg model	154
6.2.6	Errors in fidelity benchmarking	155
6.2.7	Benchmarking with evolution constraints	157
6.3	Applications	158
6.3.1	Hamiltonian parameter estimation	159
6.3.2	Target state benchmarking	162
6.3.3	Cross-platform comparisons	164
Chapter 7:	Hilbert-space ergodicity	168
7.1	Closed system dynamics	170
7.2	Open system dynamics	175
7.3	Details of experiments and simulations	178
7.3.1	Experimental details	179
7.3.2	Simulation details	180
7.4	Deriving the Erlang distribution	182
7.4.1	The probability-of-probabilities distribution	183
7.4.2	The Porter-Thomas distribution from Hilbert-space ergodicity	183
7.4.3	The Erlang distribution from Hilbert-space ergodicity	185
7.4.4	Second Hamiltonian twirling identity	187
7.4.5	Universal variance including finite-size corrections	189
7.4.6	Higher moments of subsystem bitstring probabilities	193
7.4.7	Bitstrings vs time, and time-dependence of the PoP	195
7.4.8	Failure cases	198
7.5	Generalizing to hypoexponential distributions	198
7.5.1	Porter-Thomas from random distributions	199
7.5.2	Hypoexponential distributions	200
7.5.3	Intrinsic baths	200
7.5.4	Extrinsic baths	204
7.5.5	Other noise sources	208
7.6	Effects of finite sampling	208
7.6.1	Comparing experiment and finite-sampled model PoP	210
7.6.2	Multinomial distribution and Poisson sampling	211
7.6.3	Finite-sample estimate of PoP distribution	212
7.6.4	Inverting finite sample effects	214
7.7	Cumulants of PoP distributions	217
7.7.1	Relation of hypoexponential moments to mixed state moments	218
7.7.2	From weights to Kraus operators	219
7.7.3	Entropy estimation from PoP moments	222

7.7.4 Estimator of moments	223
Coda	225

III The beyond-classically-exact regime 227

Prelude	228
Chapter 8: Benchmarking large scale simulators	231
8.1 Fidelities from approximate algorithms	232
8.2 Experiment details	237
8.2.1 Defining the experimental time unit	237
8.3 MPS cross-entropy benchmarking	239
8.4 Non-exponential fidelity decay	243
8.4.1 Fidelity response to coherent Hamiltonian errors	244
8.4.2 Non-exponential decay for Rydberg quantum simulators	247
8.5 Monte Carlo inference	251
8.5.1 Method description	252
8.5.2 Consistency checks	254
8.6 Scalability of fidelity estimation	259
8.6.1 Markovian noise	259
8.6.2 Non-Markovian noise	260
8.6.3 Sample complexity	262
Chapter 9: Quantum comparisons	264
9.1 Experimental mixed state entanglement	265
9.2 Negativity bounds and estimation	268
9.2.1 Extrapolating pure state entanglement growth	268
9.2.2 Negativity lower bounds and estimators	270
9.3 Robustness to realistic experimental errors	275
9.3.1 Error model simulations	275
9.3.2 Global coherent errors	275
9.3.3 Systematic miscalibration errors	279
9.3.4 Incoherent local errors	281
9.4 Entanglement of state preparation	288
Chapter 10: The classical cost of quantum simulation	291
10.1 Overview	291
10.2 Simulation methods	293
10.2.1 Description of the Caltech computing cluster	293
10.2.2 Schrieffer-Wolff transformations	294
10.2.3 Standard Trotterization-based TEBD-MPS	296
10.2.4 The Lightcone-MPS algorithm	297
10.2.5 Time-dependent variational principle	301
10.2.6 Noisy MPO-based simulation	302
10.2.7 Path integral formulation	303
10.2.8 Neural network based Hamiltonian simulation	305
10.3 Estimating the MPS fidelity	306
10.3.1 Extrapolating the experimental classical cost	308

10.3.2 Fitting and extrapolating the MPS fidelity	308
10.4 Negativity of truncated MPS states	310
Coda	313
Chapter 11: Outlook	315
11.1 Improving fidelities	315
11.2 Increasing system sizes	318
11.2.1 Designing for advantage	320
11.3 Exploring new physics	323
11.3.1 Quantum processing for metrology	323
11.3.2 Uncovering critical phenomena	325
11.3.3 Investigating unconventional thermalization	326
Appendix A: Experimental automated feedback	331
Appendix B: Numerical simulation methods	337
B.1 Hamiltonians and symmetries	337
B.2 Time-evolution	340
B.3 Random unitary circuits	344
Appendix C: Many-body temperature	346
C.1 Temperature	346
C.2 Infinite temperature Rydberg atom arrays	349
Bibliography	351

CHAPTER 1

Introduction

The central goal of this thesis is to address the seemingly fraught task of verifiably controlling quantum systems with high fidelity, despite explicitly attempting to do so in a regime where we both cannot fully measure the quantum state, and where no classical computer could possibly represent its complexity. To that end, we begin here by broadly describing the field of quantum science, the importance of controlling such systems, the myriad of ways researchers are attempting to do so, and finally by outlining our contributions to this continuing, monumental effort.

1.1 BACKGROUND

The theory of quantum mechanics can seem paradoxical in many ways, not least of which because though it underlies the last century of technological progress, it is strangely invisible in our day-to-day lives. Indeed, in many ways quantum mechanics is alien to our usual sensibilities and intuition. For instance, quantum particles can be in two places at once (*superposition*) and multiple such particles can be correlated so extremely that their individuality is subsumed into one collective (*entanglement*). These ‘spooky’ aspects of the theory disturbed even its originators [1] and it was not until the 1970s that experiments began conclusively demonstrating the validity and unavoidability [2] of these phenomena [3, 4]. These early efforts were recognized by the Nobel Prize in Physics in 2022 and in the decades since those pioneering discoveries, the field of experimental quantum science has flourished.

The drive to study quantum mechanics is fundamentally interesting — it has proven time and time again one of our most accurate theories of how the universe works — but much

of the modern fervor within the field is practically minded. In the early 1980s, physicist Richard Feynman [5], amongst others [6, 7], suggested that the quantum world may be too complex to simulate accurately on an ordinary, or classical, computer. Then from the 1990s continuing until today, a series of landmark results [8–13] have found classical computational problems in a wide variety of disciplines which are extremely challenging to solve on classical computers, but can be solved efficiently by carefully engineering quantum systems, or *quantum computers*; for a very exhaustive list of applications, see Ref. [14].

Importantly, these advantages are not because quantum computers are simply faster versions of ordinary computers¹, but because they operate on fundamentally different principles, which are integrally linked to the ‘spooky’ characteristics of quantum theory [15]. For instance, whereas ordinary computers are composed of *bits* which are either 0 or 1, the most basic of quantum computers are composed of *quantum bits*, or *qubits*. Qubits can also be either 0 or 1, but also any superposition of the two — you can think of a qubit as an arrow in three-dimensional space that can point up or down (1 or 0), but also forward, backward, left, right, or any other direction. We write the state of the qubit as $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$, where c_0 and c_1 are complex numbers which tell us how closely the “arrow” is pointing upwards (1) or downwards (0).

While we may represent the state of N bits with N real numbers, representing multiple qubits is markedly more complicated. For instance, the most general state of just two qubits is a superposition of all possible combinations of the individual qubits, $|\psi\rangle = c_0|00\rangle + c_1|01\rangle + c_2|10\rangle + c_3|11\rangle$, which requires us to know $2^2 = 4$ complex numbers. To describe N qubits requires 2^N complex numbers. For a system composed of only $N = 60$ qubits, the memory required for a classical computer to represent the quantum state is around ten exabytes, or roughly the entire storage capacity for all of Google. This exponential scaling underscores the difficulty of studying quantum systems with classical computers, but also represents great potential, because with only 60 qubits, we can access a state space equivalent to quintillions of classical bits. Though this alone does not guarantee the computational power of quantum computers, careful engineering and control over the state of a few qubits allows us to encode and solve the problems of interest mentioned previously.

The trouble, then, is simply to control 60 qubits.

¹In contrast, in terms of clock speed they are between a thousand to a billion times *slower*.

1.2 THE CHALLENGE OF QUANTUM SCIENCE

The possibility and promise of quantum computers to solve so many problems of broad interest has sparked a worldwide, ongoing race to develop such a system with substantial academic, governmental, and industrial investment. As of now, however, no such practically relevant device has been created. The chief reason for this comes down to a phenomena known as *quantum decoherence* [16, 17]. Without belaboring the explanation too much here — understanding decoherence fundamentally and practically will be a major thrust of this work — decoherence is essentially the loss of information about the state of a system due to its interactions with the environment.

Imagine trying to prepare an intricate meal. Working alone, you know exactly which and how many ingredients are cooked in what order, and after following the recipe, the dish will come out perfect. Now imagine, however, that whenever you open the lid of the pot to add your ingredients, a crowd of onlookers can do so as well, each adding some random spice to the pot. Very soon, the preparation will be spoiled, and the final product ruined. There is a paradoxical relationship: we want an isolated system which allows us to control it, but allows nothing else to do so. Further, the larger and longer the recipe, the more times we must open the lid, and thus the more opportunities for the dish to be ruined. Controlling quantum systems is essentially hampered by a too-many-cooks-in-the-kitchen scenario.

However, where the metaphor breaks down, and where one of the central challenges of quantum science lies, is that while the dish can be tasted and immediately deemed rotten, it is incredibly difficult to tell how, or even if, the quantum programming has gone awry. This is because measuring quantum states is *projective*. This means that though the state may be in a complex superposition of 2^N components, reading out the state of the system randomly selects one of those components, with probabilities controlled by the particular quantum program [18, 19]. Thus, quantum computers are inherently probabilistic entities, which are then further complicated by probabilistic errors. As a point of reference, while classical computers make one error every 10^{17} operations, even the best-performing current quantum devices hover around error rates of 0.1% per operation.

This difficulty was recognized since the conception of the field, spawning the vast subject of *quantum error correction* [20–23]. In these protocols, quantum information is protected from decoherence by entangling the states of many qubits, and sharing information between them such that local errors show clear signatures when measured projectively. By interspersing checks throughout the program to verify no errors have taken place — and ideally correcting them if detected — the program can thus be run error-free. However,

to realize a net gain the underlying error rates of the quantum device must be sufficiently small [24, 25] and the number of available qubits must be sufficiently large, in addition to other technical requirements such as repeated, mid-circuit readout. While there has been significant progress on these fronts, no large-scale, error-corrected system with practical ramifications has been demonstrated. Still, creating such a fully error-corrected device seems to be the only reasonable way to reliably realize the power of quantum computation.

Thus, the major goals of the field currently are generally summarized as: 1) better understand the physical errors mechanisms affecting a given quantum system, 2) measure and validate the performance of a device affected by these error sources, 3) correct or otherwise eliminate the errors. Similarly, making progress on these fronts is the overarching goal of this thesis, with several novel contributions specifically towards error characterization [26, 27], detection [26, 28], verification [29–31], and elimination [26, 32]. Though many of our results will be intentionally generally applicable to any quantum system, we will show experimental demonstrations of our techniques with a state-of-the-art array of atoms trapped in optical tweezers, for which we demonstrate enhanced scalability [33] and improved fidelities for single-qubit [28, 32, 34], two-qubit [26, 32, 34], and many-qubit operations [26, 29, 31]. Ultimately, through application of our techniques, we will demonstrate our system is controlled with such high-fidelity that it is beginning to enter the regime of *quantum advantage*, where classical computers become incapable of replicating the quantum experimental results [31].

1.3 QUANTUM PLATFORMS

Since the 1990s, many candidate platforms have been in development which are attempting to achieve the dream of useful quantum computing. While these various approaches differ in substantial, practical ways, they face the same fundamental tension between isolation and control. In this crowded landscape, a new challenger has emerged in the last decade and firmly established itself: *Rydberg atom arrays* [35, 36]. Though many of the techniques and ideas we will discuss in this thesis are explicitly general to many quantum platforms, most of our experimental demonstrations will come using a state-of-the-art Rydberg atom array. We describe details of the platform and its various error mechanisms in Chapter 2. A key facet to note, however, is that like other quantum platforms, Rydberg atom arrays can be used for performing universal gate-based quantum computation [37–40], or to study a rich assortment of spin-models [31, 41–47].

This distinction is important. There are essentially two modalities in which controlled quantum systems can be operated, either *digital* or *analog*. The same distinction can be

made classically: when trying to simulate how a particular plane will fly, for instance, one might either program the dynamical equations into a computer (digital), or build a wind tunnel¹ replicating those dynamics and test scale models (analog).

Digital quantum devices, which are sometimes distinguished as more proper quantum computers, or *quantum processors*, operate much like digital classical computers. They implement a discrete, universal gate-set, which then can be used to program any unitary circuit of interest [15, 49]. In the modern noisy, intermediate-scale quantum [50] (NISQ) era, however, compilation overhead can be inefficient [51, 52], and can increase error rates [53] for instance due to limited native expressivity [54]. Errors affecting these systems can often be abstracted into tractable forms [55–58], though this analysis is potentially fraught with oversimplification of the physical dynamics.

Analog quantum devices, also called *quantum simulators*, implement particular problems of interest by mapping them directly into the native continuous dynamics of the particular quantum system [36, 59–65]. This idea of using a quantum system to essentially “simulate itself” was the original idea of Feynman [5]. This approach largely eliminates overhead, at the cost of restricting the explorable space of problems, and making theoretical analysis in general more complex [66].

In this thesis, we explore both modalities. A major focus on this work is the particular problem of how to verify that a quantum simulator is giving trustworthy results. This task, already difficult in the digital case, is made far harder by the continuous nature of analog dynamics and errors. While historically quantum simulators have been treated as error-agnostic machines, with a concentration on low-order, supposedly-robust observables [59, 66], our goal explicitly is to validate the performance of these systems at all-orders. Our interest in digital quantum processors will mainly stem from their potential to dramatically improve the performance of quantum metrological systems, such as atomic clocks [28, 39].

1.4 OUTLINE

Now, we briefly summarize the remainder of this thesis. Overall our work is broadly split into three different parts.

In Part I we focus on systems of just a few atoms to further define the Rydberg atom array platform including its native strengths and weaknesses. In particular we will describe various errors affecting the system, and will then immediately describe how some of them may be mollified through a newly-realized error mitigation technique known as *erasure*

¹Per von Neumann, who called wind tunnels “computing devices of the so-called analogy type” [48]

conversion [67, 68]. We shall discuss how this technique enabled our record achievement of 0.999 entanglement-fidelity between two Rydberg atoms [26], in addition to a myriad of benefits such as improved quantum simulation [26], cooling to the motional ground state [32], and even enhanced loading of an optical tweezer array [33]. Finally, we will discuss how Rydberg atom arrays can be endowed with a set of universal quantum operations including arbitrary single-qubit gates [28], fixed two-qubit gates [32], and mid-circuit measurement [32]. Such operations are particularly useful for our system which can also function as precise atomic clock [69], as we demonstrate with concrete and practical methods for improving quantum metrology using digital quantum computing techniques [28].

In Part II we move to the realm of the many-body and ponder how to verify the performance of large quantum systems without incurring an exponential cost due to their exploding Hilbert space dimension. To start, we dive into the world of *random state ensembles*, and discuss newly discovered ways in which randomness emerges from natural, unitary Hamiltonian time dynamics [27, 29]. Such discoveries open new questions about the connections between quantum mechanics and statistical mechanics, with many interesting theoretical discoveries [27, 70, 71]. Exploiting the chaotic nature of this emergent randomness, we then devise protocols [29, 30] for measuring the global many-body fidelity with relatively few measurements for systems with less than roughly 30 atoms. Besides enabling benchmarking of large quantum systems, we discuss several important implications of this technique for Hamiltonian estimation and noise learning, amongst others.

In Part III we enter the beyond-classically-exact regime and face the nebulous task of benchmarking quantum systems when no classical reference is possible [31]. We will discuss why existing evaluation protocols fail because of the analog nature of errors affecting our system, but will propose and implement an extension of our many-body fidelity estimation through a many-pronged comparison against many *approximate* classical algorithms with varying degrees of entanglement. We successfully estimate the fidelity – roughly ten percent – to prepare a system of sixty atoms into a state with nearly maximal entanglement entropy. We then showcase a novel protocol for estimating the actual mixed state entanglement [72]. This allows for mutual comparison amongst state-of-the-art quantum platforms, with which we find we are competitive. Finally, we evaluate the classical cost of quantum simulation, demonstrating both the present computational power of our system, and the long-term potential of highly entangled analog quantum simulators more generally [59].

In total, the contributions of this thesis span the gamut from single atoms to many-body systems, with advancements both theoretical and experimental in the fields of quantum computing, simulation, information science, and metrology.

Part I

Rydberg atom arrays

Prelude

Understanding and minimizing the effect of errors is critical to the future of quantum science, both in the near [50] and far [24] term. Rydberg atom arrays have emerged as an exciting platform with vast scalability [73] combined with high fidelity operations [26, 39, 74, 75]; still there is much progress yet to be made. In the following chapters, we describe our contributions to this ongoing effort.

To start, we reiterate the key features of atom arrays, and of our experiment [76] specifically. We present a comprehensive summary and overview of the various major error sources affecting these systems, and explain how they may be realized in a numerical model which closely replicates experimental observations. With this model in hand, we introduce and demonstrate a new tool in the quantum toolbox, *erasure conversion* [67, 68], which we apply to improve quantum simulation [26], to cool atoms to their motional ground state [32], and even to improve system sizes (through a classical sense of erasure conversion). Through this, we demonstrate a record two-qubit entanglement fidelity with neutral atoms, with higher fidelities a question of technical improvements which we understand well.

Then we parlay these improvements to demonstrate the tools of universal quantum computation with neutral atom arrays. Specifically, we show arbitrary single-qubit gates [28], fixed two-qubit gates [32, 39], mid-circuit measurement [32], control of a potentially bosonic degree of freedom [26], and even the controllable generation of so-called hyper-entanglement [32]. With these tools we can tackle largely any computing task, but we turn our attention to metrology, and demonstrate how the marriage of an atomic clock and a quantum computer can yield improved sensing performance [28].

CHAPTER 2

Experimental setup and error modeling

In this chapter, we describe the main experimental workhorse of this research, a *Rydberg atom array*. Rydberg atom arrays have been established for nearly a decade, and as such there are various excellent resources available to learn about their basic physics. In particular, we recommend reviews in Refs. [35, 36], and for a wealth of experimental and theoretical information one can scarcely do better than consulting Refs. [76, 77].

In particular, Ref. [76] – a previous thesis from our group at Caltech – is an excellent exposition of the basics of the platform, and additionally contains comprehensive descriptions about our specific experiment. In this work, we will omit many such details, and so if ever in doubt on a technical aspect of basics of the platform, Ref. [76] is a good place to start.

2.1 TRAPPING AND CONTROLLING SINGLE ATOMS

Central to the success of any quantum platform are the qubits which compose it. In Rydberg atom arrays, the qubits of choice are single *neutral* atoms (neutral as opposed to ionic species which have been explored in detail in the study of trapped ion quantum computers [78]). The key benefits of the platform are plentiful, starting with its innate scalability. Single atoms are sourced from a magneto-optical trap of tens of millions of atoms, and loaded into a user-configurable array of tightly focused ($<1 \mu\text{m}$ waist) laser beams, known as optical tweezers¹. All atoms (qubits) are identical², meaning limitations on achievable system size originate mainly from how many optical tweezers can be created with available laser

¹The invention [79, 80] of which was recognized by the Nobel Prize in Physics in 2018.

²This is an oft repeated claim, but is really not that true – while atoms themselves are identical, true uniformity requires their environment, including their trapping conditions, be uniform as well. Such metrics are typically hard to make better than a $\approx 0.1\%$ homogeneity level [81]

power: experiments with several hundred [42, 43] or even a few thousand such traps [82–84] are readily achievable, and recently a landmark, state-of-the-art demonstration of a twelve thousand trap array filled with over six thousand atoms was demonstrated [73].

Still, having many atoms counts for very little if the fidelity of controlling them is poor. Many factors affect the fidelity of the platform, but the most basic is loss of atoms from their traps. Such lossy type errors are a subset of the broader class of *leakage* errors. In order to minimize such effects, atoms are trapped in optical tweezers in a vacuum chamber at a pressure of $\approx 10^{-12}$ Torr, or about the atmospheric pressure on the surface of the Moon (a billionth of a millionth of Earth’s atmosphere). This results in few collisional interactions with background gas molecules, leading to innate trapping lifetimes on the order of tens of minutes [73, 85] (even hours in cryogenic environments [86]).

With atoms firmly held in place, the internal electronic state of the atoms can then be manipulated with microwave or laser sources to perform single qubit gates, again with high fidelity (>0.999). The state of the atoms can be detected by fluorescence imaging while atoms remain trapped, and high-resolution microscopy techniques pioneered over the last twenty years [87, 88] enable fully resolving atoms spaced by a few micron. Further, the optical tweezer array geometry is programmable in near real-time, meaning atoms can be rearranged into largely any configuration desired in one [81], two [43, 89, 90], or even three [91, 92] dimensions.

In order to generate entanglement, we must make multiple qubits interact, which we accomplish through excitation to high-energy *Rydberg states*. When in the Rydberg state, atoms experience a long-range, distance-dependent interaction based on a second-order induced dipole effect. In the early 2000s, this interaction was recognized as a potent resource due to the so-called *Rydberg blockade* mechanism [93, 94]. As the simplest example, when two atoms are close enough together (a few micron separation), excitation of one to the Rydberg state is allowed, but the energy of the doubly-excited state can be shifted dramatically. This precludes occupation of such doubly-excited states, meaning when the two atoms are simultaneously driven to the Rydberg state the result is a single Rydberg excitation shared across the two atoms, i.e. an antiferromagnetic Bell state; we will explore this process in great detail in Chapter 3. When operating in a digital modus, this enables a variety of schemes for performing high-fidelity controlled-phase gates [39, 74, 75, 95] (as we will explore and discuss in Chapter 4).

2.1.1 The Rydberg Hamiltonian

The Rydberg interaction has further ramifications for systems with more than two atoms. In this work, we will almost always consider a one-dimensional chain of atoms with ideally uniform spacing. As such, the Rydberg Hamiltonian is well approximated by

$$\hat{H}/\hbar = \Omega \sum_i \hat{S}_i^x - \Delta \sum_i \hat{n}_i + \frac{C_6}{a^6} \sum_{i>j} \frac{\hat{n}_i \hat{n}_j}{|i-j|^6} \quad (2.1)$$

describing a set of interacting two-level systems, labeled by site indices i and j , driven by a laser with Rabi frequency Ω and detuning Δ , with \hbar being Planck's constant. The interaction strength is determined by the so-called C_6 coefficient and the lattice spacing a . Operators are $\hat{S}_i^x = (|1\rangle_i \langle 0|_i + |0\rangle_i \langle 1|_i)/2$ and $\hat{n}_i = |1\rangle_i \langle 1|_i$, where $|0\rangle_i$ and $|1\rangle_i$ denote the electronic ground and Rydberg states at site i , respectively.

The Rydberg blockade mechanism is still of central importance in this case. For typical parameters, the nearest neighbor interaction strength (C_6/a^6) is a few hundred MHz, as compared to typical achievable Rabi frequencies of a few MHz. As such, simultaneous excitation of neighboring atoms to Rydberg states are strongly suppressed. To first order the blockade reduces the Hilbert space dimension from 2^N to $\text{Fib}(N+2) \approx 1.62^N$, where Fib is the Fibonacci function (see Appendix C). This means that for a system size of N atoms, the effective Hilbert space size is only that of $\sim 0.7 N$ qubits. In this case the Hilbert space is complete [96], but becomes non-factorizable, i.e. if you bipartition a system of Rydberg atoms with Hilbert space dimension $\dim(\mathcal{H})$ into two halves, left and right with Hilbert space dimensions $\dim(\mathcal{H}_L)$ and $\dim(\mathcal{H}_R)$, respectively, then $\dim(\mathcal{H}_L) \times \dim(\mathcal{H}_R) \neq \dim(\mathcal{H})$. This has practical ramifications for the growth of entanglement in these systems, while also enabling efficient classical simulation methods in some cases [97, 98].

Often, we will parameterize the time-axis during evolution under the Rydberg Hamiltonian in terms of *cycles*, given by $t_{\text{cycle}} = 2\pi/\Omega$. This timescale is chosen as Ω is the dominant energy-scale in the system, after the blockading nearest neighbor interaction energy. As such, entanglement growth in the blockaded Hilbert space is primarily constrained [31] by t_{cycle} .

2.2 STRONTIUM

Through the many iterations of the Rydberg atom array platform, many atom species have been utilized. At the outset, it is important to note that as of now, there is no clear consensus as to which element is 'best,' and in fact it may be that none can actually hold the crown.

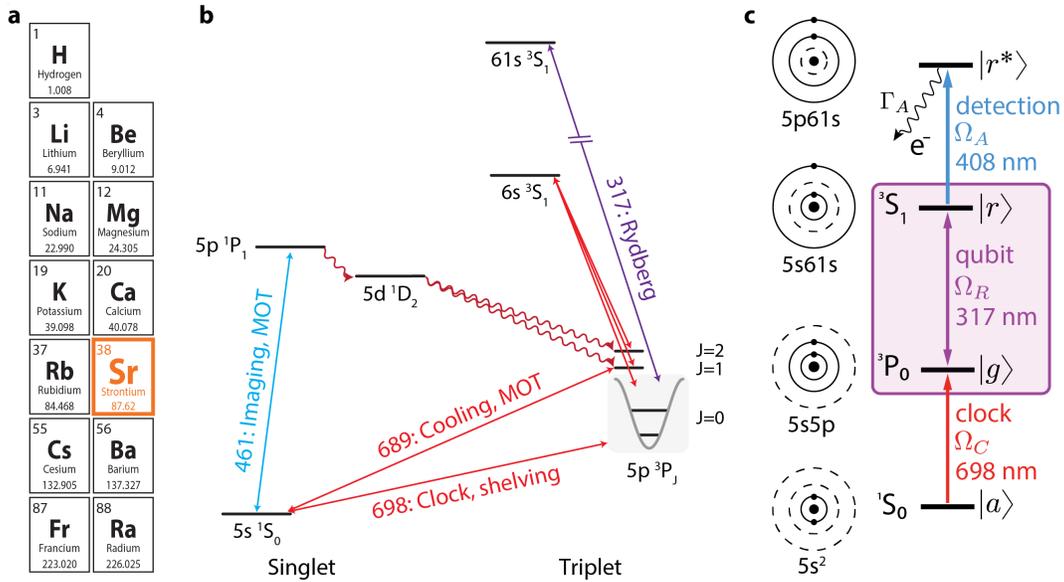


Figure 2.1: **Relevant level structure of strontium a**, Strontium in the periodic table. **b,c**, Simplified level structures of strontium with selected transitions labeled along with their wavelength, highlighting states important to this work of this thesis; for a more complete version, see for instance Ref. [107]. In general, for many purposes of this thesis we will be concerned with an effective two-level qubit subspace defined between the $5s5p\ ^3P_0$ clock state and the $5s61s\ ^3S_1$ Rydberg state. Additional levels are for instance used for atom trapping, imaging, and cooling, as well as detection of the Rydberg state through autoionization [34]. As the exception, in Chapter 4 we will focus on the qubit defined between $5s^2\ ^1S_0$ and the clock state, in order to realize a universal digital quantum processor with strontium atoms. We operate in the sideband resolved regime for the ultranarrow clock transition, which allows us to address individual motional transitions. We will use this ability to great effect for atom cooling as well as mid-circuit measurement in Chapter 4.

Each choice of element comes with pros and cons, both to the complexity of control required, and the richness of physics attainable.

Early experiments utilized alkali atoms which are the historical workhorse of the atomic physics community due to their largely hydrogenic level structure. Such alkali-based experiment are still in use in modern, state-of-the-art experiments, but in recent years a second-generation of experiments has been developed which utilize alkaline-earth(-like) atoms such as strontium [99–102] or ytterbium [103–106]. These elements have two valence electrons, which leads to a somewhat more complex level structure, but which enables a variety of striking applications. Ours is one of these original experiments [99] (alongside Refs. [100, 103]), and utilizes the bosonic strontium-88 species.

Strontium has a number of enticing advantages. Its two valence electrons lead to a level structure composed of both singlet and triplet electronic states; the level structure of states

relevant to this work are shown in Fig. 2.1. The absolute ground state is the singlet state $5s^2\ ^1S_0$. Within the singlet manifold, the only transition of major importance for our purposes is the broad transition to $5s5p\ ^1P_1$ which is used for MOT cooling, as well as atomic imaging [76].

The more interesting transitions are those between the singlet and triplet manifolds, which are only weakly-allowed due to spin-orbit interactions, leading to narrow linewidths. For instance, the transition between $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_1$ is only 7 kHz broad (much narrower than typical trap frequencies of ≈ 70 kHz, and as such can be used for atomic cooling schemes such as sideband [99] or Sisyphus [85] cooling to near the motional ground state, in addition to providing a second stage of MOT cooling [76, 99].

Perhaps the most interesting transition in strontium is that from $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_0$. This transition is in fact double-forbidden (both due to the required spin-flip, as well as by the apparent violation of angular momentum conservation), and for bosonic isotopes like strontium-88 is approximately [108] enabled only through application of a strong magnetic field to effectively admix 3P_1 and 3P_0 [109]. Once prepared in 3P_0 , the atom has an effectively infinite lifetime compared to typical experimental timescales, though in reality is limited to around 1 s due to technical effects like Raman scattering from the trapping light. It is for this reason that the transition to 3P_0 is the basis for the current most precise atomic clocks in the world [110, 111], which use this ultra-narrow transition (employing fermionic isotopes rather than bosonic) to stabilize the frequency of lasers locked to state-of-the-art high finesse cavities. Thus, 3P_0 is often called the *clock state*, and the $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_0$ transition is similarly called the *clock transition*; we will employ this parlance throughout this work.

Strontium has thus been studied in a metrological context for several decades. In fact, one of the earliest works of our experiment [69] (alongside Ref. [112]) was to demonstrate running such an optical atomic clock using a tweezer array with various benefits such as single atom resolution and reduced deadtime. In Chapter 4, we will extensively discuss utilizing the clock transition to enable digital quantum computation schemes and to realize enhanced metrological protocols. However, for much of the rest of this work, we will disregard this myriad of interesting physics, and will instead solely employ the clock state because of its long-lived, metastable nature.

The long lifetime of the clock state is important because when making atoms interact through the Rydberg interaction, we ideally want to excite them to isotropic Rydberg S orbitals. However, given that the absolute ground state of our system is also such an S state, this cannot be accomplished with direct photon excitation (and even if it could, such

excitations would be ≈ 218 nm, a very inhospitable wavelength to work with). Thus, a typical experimental sequence will begin by first transferring all atoms to the clock state, which is then treated as a metastable ground state from which single photon Rydberg excitation can take place. We note that this is fundamentally different than the typical approach of most alkali atom arrays, which do not have access to such a long-lived, metastable state, and so instead generally rely on two-photon transitions to enable Rydberg interactions. We note that the long lifetime of the clock state will additionally be integral to our implementation of error mitigation techniques known as *erasure conversion*, as shall be discussed in Chapter 3.

Besides preparation, alkaline-earth atoms also feature a unique benefit in terms of Rydberg state readout. This is because only one of the two valence electrons will actually be excited to the high-lying Rydberg orbital, leaving the other valence electron in the low lying manifold. However, this secondary electron can also be excited to an intermediate energy state which quickly leads to electron-electron interactions which ionize the atom, termed *autoionization* (Fig. 2.1c). The resulting ion is dark to the normal fluorescence imaging transitions, and is typically quickly expelled from the optical tweezer due to stray electric fields. In contrast, our normal fluorescence imaging detects atoms in the clock state with high fidelity. Thus, we can perform state discrimination between the clock and Rydberg states by first performing Rydberg excitation, then applying an autoionizing pulse, and finally imaging. We demonstrated this effect in Ref. [34], and showed it could reach discrimination fidelities in excess of 0.999. Further, the autoionizing transition is particularly well adapted for performing single site control on Rydberg excitations via targeted light shifts [113], though we do not use this capability in this work.

To summarize, in a typical experimental sequence atoms are first loaded in the $5s^2\ ^1S_0$ state, and then cooled on the narrow-line $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^3P_1$ (689 nm) transition close to their motional ground state. When studying analog quantum evolution in the clock-Rydberg manifold, atoms are then initialized to the $5s5p\ ^3P_0$ (698 nm) clock state through a combination of coherent drive and incoherent pumping. We treat the clock state as a metastable qubit ground state, $|0\rangle$, and subsequently drive to the $5s61s\ ^3S_1, m_J=0$ (317 nm) Rydberg state, $|1\rangle$. Following Hamiltonian evolution, state readout is performed using the autoionizing transition $5s61s\ ^3S_1, m_J=0 \leftrightarrow 5p_{3/2}61s_{1/2}$ (408 nm, $J=1, m_J=\pm 1$) which rapidly ionizes atoms in the Rydberg state leaving them dark to our fluorescence imaging. Atoms in the clock state are pumped into the imaging cycle, allowing us to measure qubit state with atomic fluorescence measurements on the $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^1P_1$ (461 nm) transition. The result of each measurement is a *bitstring* for N atoms, $z \in \{0, 1\}^{\otimes N}$, where e.g. $z = 100101001$. We can repeat measurements many times to accrue many such

bitstrings, from which we make inferences and observations about the many-body physics we wish to study.

2.3 ERRORS FOR RYDBERG ATOM ARRAYS

The Rydberg Hamiltonian (Eq. (2.1)) is the paradigmatic model for our experiment. However, in practice quantum systems are not perfectly isolated and can undesirably interact with their environment. This could either be due to control parameters in the Hamiltonian drifting over time (e.g. as the laser intensity randomly fluctuates) or coupling to extrinsic quantum degrees of freedom (e.g. spontaneous decay of excited states).

Much of the experimental work of this thesis is focused on the problem of identifying the signatures of these errors using experimental data, for instance by measuring the fidelity of performing a given quantum operation, but it is equally important to develop a well-constructed error model to understand the impact of errors in realistic scenarios. We set out to define such an error model here, which we will reference and use extensively throughout the remainder of this work. After incorporating all known error effects, we generally find our error model agrees very well with experiment even on the level of many-body bitstring probability distributions, for instance typified by studying the difference between probability distributions during quench evolution of the experiment versus either error-free or error-model simulation (Fig. 2.2). While in this Section we take a descriptivist approach to the matter of noise modelling, in Chapter 7 we approach it both theoretically and experimentally from the perspective of learning noise sources directly from their effects on quantum many-body observables.

At the outset, it is important to note that there are many ways of simulating open system dynamics in quantum systems, for instance through master equation-based approaches which explicitly consider time evolution of the experimental mixed state under appropriate jump operators. While such an approach is effective for simple and small systems, it requires manipulating the 4^N elements of the density matrix which is prohibitive for larger systems. We wish to use our error model for large systems (> 20 atoms), and include the entire gamut of complex error sources. As such, we forgo the master equation, and instead perform our open system dynamics simulations using a Monte Carlo wavefunction [114] approach.

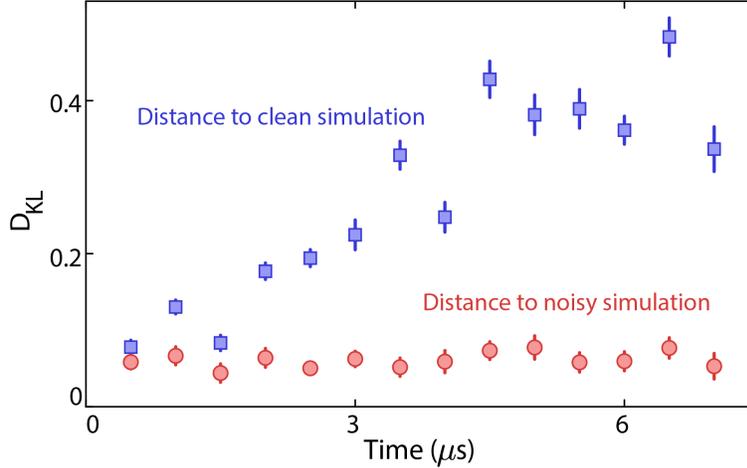


Figure 2.2: **Quantitative evaluation of *ab initio* error model:** Kullback-Leibler (KL) divergence of bitstring probability distributions between the ten-qubit Rydberg experiment and an error-free model (blue square) and an *ab initio* error model (red circle). When the experimental data is benchmarked against an error-free model, the KL divergence increases over time since the presence of decoherence effects and control imperfections are not taken into account in the error-free model. In the case of the error model, however, we find a significant reduction in the KL divergence between the experiment and error model at *all* times.

2.3.1 Summary

As shown in Eq. (2.1), the Rydberg Hamiltonian is well approximated by

$$H = \Omega \sum_i S_i^x - \Delta \sum_i n_i + C_6 \sum_{i>j} \frac{n_i n_j}{|\vec{x}^{(i)} - \vec{x}^{(j)}|^6},$$

where the interatomic spacing is written as $|\vec{x}_i - \vec{x}_j|$.

To a good approximation, error sources in the experiment can be written via the following transformations (where notations are defined in subsequent sections, e.g. \mathbb{LP} is a low-pass function given by an AOM bandwidth)

$$\Omega \rightarrow \Omega_0 + \Omega_c + \Omega_{\text{dc}}(\sigma_I) + \Omega_{\text{ac}}(\omega_I, t) + \Omega_{\text{pt}}(\sigma_p) + \Omega_{\text{bs}}^{(i)}(T) \quad (2.2)$$

$$\Omega \rightarrow \mathbb{LP}[\Omega] \quad (2.3)$$

$$\Delta \rightarrow \Delta_0 + \Delta_c + \Delta_{\text{dc}}(\sigma_\nu) + \Delta_{\text{ac}}(\omega_\nu, t) + \Delta_{\text{dop}}^{(i)}(T) \quad (2.4)$$

$$\Delta \rightarrow \mathbb{LP}[\Delta] + \kappa\Omega \quad (2.5)$$

$$\vec{x}_i \rightarrow a_0 i + \delta^{(i)}(T) \quad (2.6)$$

In the following, we shall explain each of these terms, categorizing them largely by the source of the error. In practice, we typically account for all such errors, though the relevance of any

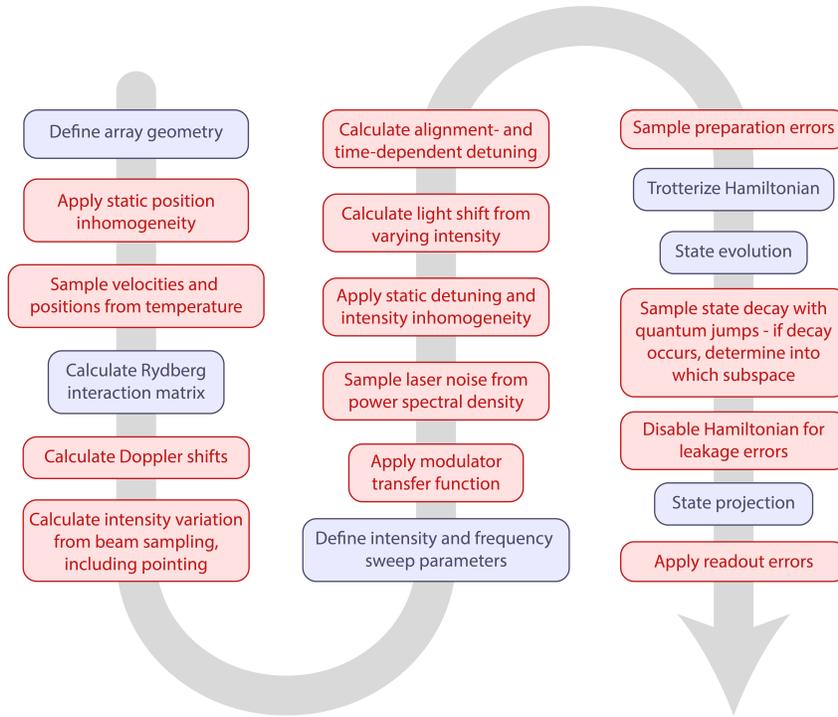


Figure 2.3: **Errors in Rydberg atom arrays.** Various physical error sources affecting the Rydberg atom array pipeline. Here we list the general structure and components of our error model, and throughout the text we will describe each error source, including ways in which to calibrate and mollify their effects. Above, blue boxes are error-free operations, while red boxes denote possible error cases

given error can vary dramatically with the type of experiment being performed (Fig. 11.1). A general flowchart for how such errors are captured in a typical sequence is shown in Fig. 2.3, and a summary of major error sources and general information about them is shown in Table 2.1. Further, we note that several, but not all, of these errors are covered in one way or another in Ref. [115].

First, a few important notes.

1) Throughout we will often distinguish between “correlated” and “uncorrelated” noise sources. By *correlated*, we generally refer to spatial correlation, i.e. a noise source which affects all atoms in a joint manner. While this is not the same as a high-weight error (i.e. a non-factorizable error acting on multiple qubits), it does lead to spatial correlations in the experimental mixed state.

2) Our error model includes numerous *leakage* errors, errors which leave the qubit subspace, such as decay to a different manifold or imperfect state preparation. To account for such errors, we could assume a higher dimensional particle subspace like a qutrit or higher.

Noise channel	Spatially-correlated?	Time-dependent?	Intensity	Frequency
DC intensity	✓	×	✓	×
AC intensity	✓	✓	✓	×
DC frequency	✓	×	×	✓
AC frequency	✓	✓	×	✓
Poynting	✓	×	✓	×
Doppler shifts	×	×	×	✓
Beam sampling	×	×	✓	×
Interaction variation	×	×	-	-
Spontaneous decay	×	✓	-	-

Table 2.1: **Error sources.** Note that in general most “intensity” noise sources also can impart frequency noise if the system experiences a self-light-shift. Given that most error source have a unique set of truth values, they can be isolated and measured via *in situ* atomic signals through measurement of on-site observables and connected correlation functions.

However, doing so would become quite computationally expensive. Instead, we take a semi-classical approach. We still treat each atom as a qubit, but additionally store a flag array (linear in system size) stating if the qubit has leaked or not. If the qubit does leak, we make it inert, turning off all associated Hamiltonian terms¹. The quantum state vector is then not a trustworthy reference, and so cannot be used for calculating state overlaps etc. Still, we can still incorporate such effects into measurements of many-body fidelity, etc. via the flag vector, and the simulation results are still accurate on the level of bitstring samples.

3) We will typically make a frozen-core type approximation, whereby many error sources are assumed to be static on the timescale of typical experiments. We will try to point out such cases.

4) As a combination of points (1) and (3) above, many of the noise sources we consider are explicitly non-Markovian and in some cases are high-weight. Such errors can have dramatic effects not only on the fidelity of performing quantum operations, but can also non-trivially change how fidelity decays during erroneous evolution, a matter we shall discuss at length in Chapter 8.

¹We make the explicit assumption that atoms which leak out of the qubit subspace cannot reenter, for which we see no evidence to the contrary on the level of observables we study.

2.3.2 Laser errors

Some of the most significant error sources affecting the neutral atom quantum platform arise from imperfect control of lasers which are used for performing atomic excitation. In general (and simplifying somewhat), a laser’s output has a stochastically varying amplitude and frequency, both of which can be largely captured through a power spectral density (PSD). Here, we break this PSD up into two basic parts: DC or “low frequencies,” and AC or “high frequencies.” While there is no firm definition of what makes for DC noise or AC noise, we generally consider AC noise to arise from the the portion of the PSD with frequencies higher than a typical experimental duty cycle, while the DC noise arises from the portion of the PSD with frequencies lower than the typical experimental duty cycle. Thus, we often will describe DC noise as “shot-to-shot” because it is roughly constant throughout a single experimental realization, but can vary from one such realization to the next. As a general note, the laser is generally applied equally across the entire array, thus making such errors correlated.

Frequency noise

To best understand frequency noise arising from the laser, we consider its two components (DC and AC) separately.

DC frequency noise The DC frequency noise is symbolized by $\Delta_{\text{dc}}(\sigma_\nu)$. There are multiple ways to estimate this contribution, but we find the simplest, and likely most accurate is to directly use an atomic signal. In particular, we import techniques from atomic clocks [69] to lock the frequency of the Rydberg transition directly to the atoms. Using as low a Rabi frequency as possible (to narrow the transition’s spectral width), we perform two-point locking [69], feeding back to the frequency of an AOM driving the transition (Fig. 2.4a). By recording the locked error signals over time, we have very fine resolution of the shot-to-shot transition frequency variations. While these variations are a convolution of both laser frequency noise and thermal noise, the laser frequency noise can be extracted either by assuming the thermal noise from independent measurements (discussed below), or by looking at connected correlation functions of the in-lock signal. Connected correlation functions subtract off the effect of uncorrelated noise sources, like thermal noise, leaving only correlated laser noise. We typically find the DC frequency variation is largely Gaussian, with a standard deviation of σ_ν , from which we sample for our *ab initio* error model. The DC frequency noise can also be cross-checked via comparison to sensitive atomic signals like Ramsey spectroscopy.

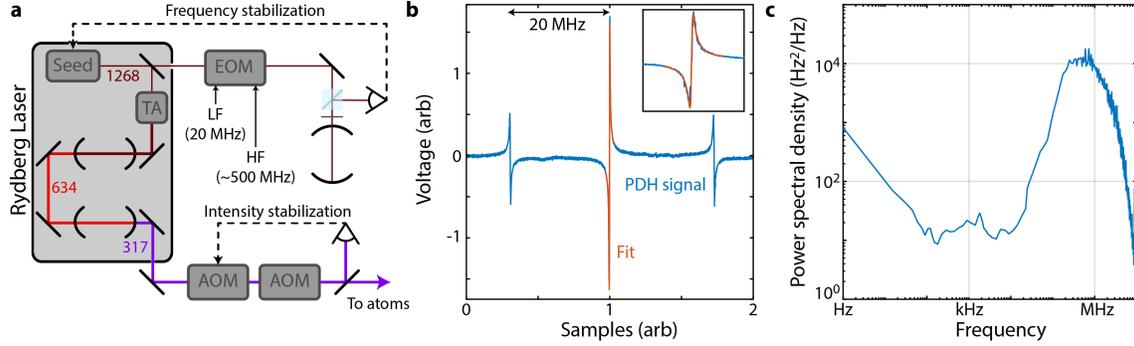


Figure 2.4: **Measuring laser phase noise.** **a**, Simplified setup of our Rydberg laser, featuring a seed laser which is then amplified via tapered amplifier (TA) before passing through two doubling cavities. The resulting power is ~ 300 mW (though with a more recent upgrade to an RFA system, we achieve up to 3 W). Also shown are the PDH locking setup for stabilizing the laser frequency, and the intensity stabilization setup post-laser. **b**, We perform PDH locking to a high-finesse (finesse of 14000, linewidth of 110 kHz) cavity. Modulation is provided by an EOM which generates sidebands at both low and high frequencies. High frequency sidebands allow us to tune the locked frequency across the entire cavity FSR. The low frequency provides the PDH locking modulation, while also providing a stable frequency reference for converting from voltage to frequency. **c**, We take long voltage time traces of the in-lock PDH signal, which we then convert to frequency via the slope calibration in **b**. Then, we convert to a power spectral density which (accounting for cavity-roll-off factors, see text), allows us to generate frequency noise traces for our system.

AC frequency noise To account for time-dependent (AC) frequency noise ($\Delta_{ac}(\omega_\nu, t)$), we must have some estimate of the power spectral density (PSD), ω_ν , of the noise.

Our approach takes advantage of the fact that we lock the Rydberg laser to a stable frequency reference, a high-finesse cavity (Fig. 2.4a). To lock to the cavity, light is modulated with an EOM with two tones, a low-frequency (LF) tone at 20 MHz, and a high-frequency (HF) tone ranging from 0 to 1.5 GHz. The HF tone allows us to span the entire free spectral range of the cavity, while the LF allows us to lock the laser with Pound-Drever-Hall (PDH) modulation [116]. In the unlocked state, the cavity rejection signal shows a characteristic trimodal structure (Fig. 2.4b), with features separated by the LF modulation frequency. By fitting a linear slope to the central PDH feature (Fig. 2.4b, inset) this provides an unambiguous conversion from voltage (as measured on an oscilloscope) to frequency deviations of the laser.

We can then lock the laser, measure the continuous-time voltage, and convert it into frequency. By taking the Fourier transform, we then arrive at the PSD (Fig. 2.4c). The PSD can then be sampled to generate unlimited time-dependent frequency noise traces. See

Ref. [117] for discussion of the technical specifics of this process.

However, first we multiply by the cavity roll-off factor [117, 118], as the error signal underestimates high frequency noise because the cavity acts as an efficient low-pass filter. The cavity roll-off factor is given by

$$\mathcal{C}(f) = \frac{1}{\mathcal{F}} + \frac{1 - \frac{1}{\mathcal{F}}}{1 + 4(f/\Gamma)^2}, \quad (2.7)$$

where f is the frequency, $\mathcal{F} = 14000$ is the finesse and $\Gamma = 110$ kHz is the cavity linewidth. In our particular circumstance, we need to perform further conditioning, given that we measure frequency noise on the infrared laser light, which is then quadrupled through two doubling cavities. This quadrupling up-converts the frequency noise, while the doubling cavities act as effective low-pass filters on the light (which also converts frequency noise into intensity noise).

Thus, the effective cavity roll-off factor is given by $\mathcal{C}_{\text{eff}} = \mathcal{C}/(\mathcal{C}_{\text{FHG}}\mathcal{C}_{\text{SHG}})$, where \mathcal{C}_{FHG} is the roll-off for the fourth-harmonic generation (FHG) cavity, and similarly for the second-harmonic generation cavity (SHG). For our particular laser, we find the resulting PSD is largely peaked around 1 MHz, which we attribute to the servo bump of the PDH lock.

However, even with all this care, it is still entirely possible that the resulting PSD does not correctly match the frequency noise seen by the atoms. There are a number of elements between the laser PDH lock and the atoms themselves, and so this PSD should be considered heuristic at best; still we find that using it matches sensitive experimental observables quite well. We note that there are means of directly measuring the frequency PSD through so-called spin-locking methods [39], which can help to confirm the findings here.

Intensity noise

As with frequency noise, we separate the contributions for intensity noise into slow, DC intensity variations, and fast, AC variations. However, unlike the case of frequency noise, we measure both fast and slow noise directly on a fast photodiode. In particular, we pick off a small fraction of the light just before it reaches the atoms (Fig. 2.4a) and direct it onto a fast avalanche photodiode (50 MHz bandwidth), being careful to make sure the measurement is not limited by the intrinsic noise of the sensor. The voltage measured is directly proportional to the laser intensity, up to small offsets which can be subtracted off.

There are multiple ways in which we could analyze the resulting voltage signal, for instance the most naïve would be to simply take a very long time trace with the laser and AOMs continuously on, perform the Fourier transform, and acquire the PSD down to very low

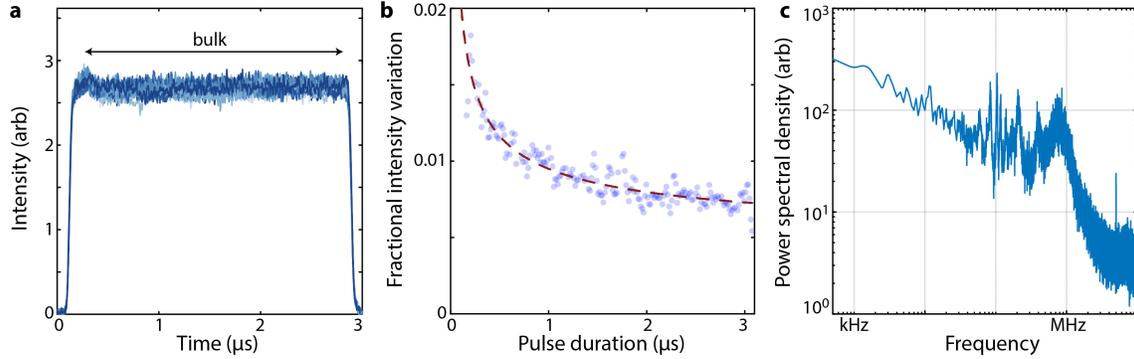


Figure 2.5: **Characteristics of intensity noise.** **a**, We measure intensity noise traces on a photodiode (schematic shown in Fig. 2.4). Each pulse (recorded at the normal experimental duty cycle), is characterized by turn-on and -off dynamics, as well as a bulk period. **b**, We study fluctuations in the bulk by fractional intensity variation, i.e. the standard deviation of the pulse area over many trials at a given pulse length, T , divided by the average pulse area. The result decreases nearly as $1/\sqrt{T}$, the expectation for white noise, before settling to a value around ~ 0.007 . We stitch multiple bulk measurements together (subtracting off the mean for each) and take the Fourier transform to acquire the intensity noise power spectral density. The result is largely similar to that found from taking the Fourier transform of a single long time trace (not shown), but in principle is more accurate to the AOM thermalization dynamics at the experimental duty cycle.

frequencies. However, with some testing, we find this is not the best approach to find the actual Rabi frequency at the atoms, because the noise spectra of a very long continuous beam of light can be different than short pulses of light with a very low duty cycle, and the latter is more accurate to the experiment. Thus we measure the intensity noise by running the experiment at its normal duty cycle, and recording pulses, $I^{(i)}(t)$, of varying length on the photodiode (Fig. 2.5a).

DC intensity noise To find the DC intensity noise, $\Omega_{\text{dc}}(\sigma_I)$, we measure the $I^{(i)}(t)$ for various total pulse durations, T , and then calculate the fractional intensity uncertainty, σ_I over iterations, i . We find that the result decays nearly as $1/\sqrt{T}$, the expectation for white noise, before settling to a value around ~ 0.007 (Fig. 2.5b). We further confirm this value via auxiliary measurements on cameras and other photodiodes (not shown). We take this long-time value as the DC intensity variation. To calculate the DC Rabi frequency we approximate that $\sigma_{\Omega} \approx \sigma_I/2$, which is true for small variations.

AC intensity noise To find the AC intensity noise, $\Omega_{\text{ac}}(\omega_I, t)$, we consider the raw intensity traces $I^{(i)}(t)$, and truncate them to only the bulk of the pulse (i.e. without turn-on and -off

dynamics). We then stitch these fragments together, subtracting off the mean of each, and take the Fourier transform to arrive at the intensity noise PSD (Fig. 2.5c). We note the result is largely similar to the PSD found from just a single very long intensity noise trace on the photodiode, but is in principal more accurate to the exact experimental thermalization. For our particular system we find the PSD peaks around a few hundred kHz, which we attribute to the servo bump of the intensity stabilization system. However, it is also possible that some of the AC frequency noise is converted into intensity noise via the doubling cavities we employ downstream from the frequency stabilization setup (Fig. 2.4a).

AOM dynamics

In order to address the atoms, the laser beam is passed through a series of acousto-optic modulators (AOMs) which control its intensity and frequency¹ (Fig. 2.4a). However, such devices cannot modulate the laser infinitely quickly, even in order to turn it on or off. For our particular AOMs, and our particular beam waist inside the AOM, we find a characteristic response time of ≈ 35 ns or ≈ 14 MHz bandwidth. Thus, any time-dependent variations in laser intensity or frequency are essentially modulated by a low-pass filter with that cutoff frequency (which we symbolize as $\mathbb{LP}[\Omega]$ and $\mathbb{LP}[\Delta]$). Typically, we do not explicitly account for this effect in most error model simulations, but note that it can non-trivially alter programmed pulse areas, and thus acts essentially as a form of miscalibration error.

There is a far more pernicious problem associated with AOM response dynamics than just their slow bandwidth, though. As a brief review, AOMs essentially work by passing an acoustic wave through a crystal, changing its refractive index with a periodic modulation. This periodic modulation then acts as a grating to incoming light, diffracting it and altering its frequency slightly. An alternative picture is to consider the acoustic wave as a propagating phonon mode, which then scatters off an incoming photon. For our purposes, the wave-based picture is more helpful.

When an RF signal is sent to the AOM, it is transduced to the acoustic wave via a piezoelectric element at one end of the crystal. The wave then propagates across the crystal at the finite speed of sound. However, impinging laser beams have a finite extent, and thus there is a finite time required for the acoustic wave to traverse the beam profile. At very early times, when the wave has not crossed the full profile, the frequency will be Fourier broadened significantly, which will then sharpen over time as the acoustic wave crosses the entire

¹We use two AOMs in order to vary the frequency while keeping the beam position constant. Such a task is often accomplished with a double pass AOM setup, but UV AOMs are very sensitive to polarization, and thus standard double pass methods are not immediately applicable. It is possible to do double passing via the use of prisms, or carefully aligned lens, with one such example shown in Ref. [119].

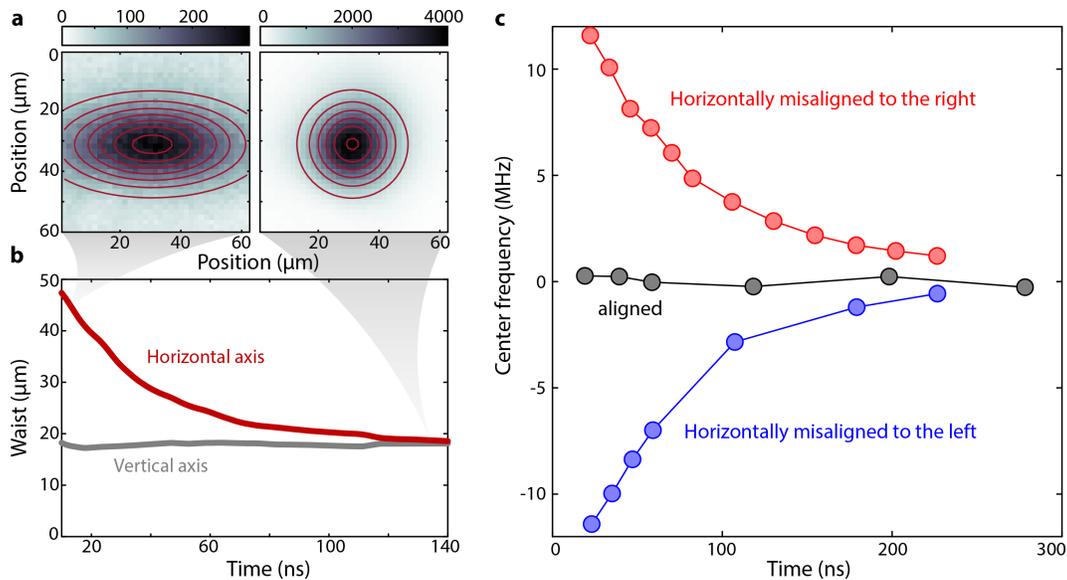


Figure 2.6: **Finite response time effects of the AOM.** **a,b**, The AOM for modulating the laser intensity and frequency (see Fig 2.4a) has a finite response bandwidth due to the time needed for the acoustic wave in the crystal to traverse the beam extent. This means at early times the beam is significantly broadened along the modulation axis (though not along the orthogonal axis), introducing a time-dependent, power-law intensity if the laser is not perfectly aligned to the atoms. **c**. Because the angle of diffraction from the AOM is proportional to the frequency imparted, the effect in **b** also leads to a time-dependent *frequency* if the beam is misaligned. For three different choices of alignment (well-aligned, misaligned to the right, misaligned to the left) we show the experimentally measured center frequency of the Rydberg excitation as a function of time. The misaligned cases converge as a power-law towards the well-aligned case at late times. Accounting for these effects thus necessitates careful alignment, which we discuss in Appendix A.

beam. What this means is that at early times after turning on the AOM, the outgoing beam will be broadened.

We observe this effect experimentally. Using a camera in an atom-equivalent focal plane of the laser after the AOM, we record the beam profile as a function of pulse duration (Fig. 2.6ab). We find that along the vertical axis (which is orthogonal to the AOM modulation), the beam profile has a constant fitted waist at all times. However, for the horizontal axis (which is parallel to the AOM modulation), the beam waist is greatly broadened at early times and decays with a power law form over the course of over a hundred nanoseconds. This means that if the beam is not well aligned to the atoms, the atoms will experience a time-dependent intensity as the Fourier components of the modulation sharpen.

However, intensity variations over time are not the only problem; there are also *frequency* variations over time. This is because the angle of diffraction from the AOM is proportional to

the frequency modulation of the light, so the broad beams at early times also have a frequency gradient across their entire horizontal span. This effect we also resolve experimentally (Fig. 2.6c). We intentionally misalign the atoms along the horizontal axis with respect to the beam center, and measure the corresponding center frequency. We see a clear power law convergence over time, much the same as for the beam waist in Fig. 2.6b, which narrows towards the well aligned value at late times.

To sum up, the AOM introduces a time-dependent modulation due to the finite speed of sound of the crystal, which leads to both time-dependent intensity and frequency variations if the atom is misaligned by as much as a micron. Thus it is extremely important to implement a robust alignment procedure (which we discuss in Appendix A) to waylay this effect. We note that alternatively one might consider using a fiber coupled system which would mix the spatially-dependent frequencies and intensities, thus transforming the well-defined detuning into a time-dependent frequency noise which might be preferred for some applications. Fibers are difficult to make work at high power for UV lasers, though there is some work in this regard [119–122].

Other effects

Various other laser-based effects can harm the fidelity of quantum operations, including the intensity-dependent detuning, beam drifts, and calibration errors.

Light shifts The laser detuning experienced by the atoms is weakly dependent on the laser intensity itself [34], and thus any variations in Ω – either programmed or erroneous – will produce a corresponding error in Δ , with scaling coefficient κ . This effect is often called a *self-light-shift*.

Poynting noise The laser is generally well-aligned to the atoms, but its exact position (and angle) still fluctuates on a shot-to-shot basis. This is known as Poynting instability¹. We measure such variations on a camera placed in a atom-equivalent focal plane of the laser, and find variations are on the order of a few hundred nanometer, as compared to beam waist-radii of a few tens of micron. Besides this variation, we also generally assume a static offset of around 200 nanometer due to imperfect alignment. Thus the effect is small, though still accounted for, and it also can have some amplifying effects on temperature-based errors, as we shall discuss.

¹I still have never had a satisfactory answer as to whether this should be “pointing”, like where the laser is pointing, or “Poynting”, like the energy flux vector. Given the ambiguity, I choose to use Poynting. However, during the course of the revision of this thesis it was pointed out by Nick Hutzler that a perfectly stable beam would still have “Poynting” fluctuations if its intensity varied, and thus maybe “pointing” should be preferred.

Calibration errors While we spend great pains to consistently and accurately feedback on the Hamiltonian parameters (see Appendix A), both the Rabi frequency (intensity) and detuning (frequency) of the laser are subject to miscalibration control errors, denoted as Ω_c and Δ_c . For many applications like many-body fidelity benchmarking such effects are somewhat minor, given that we can always accurately measure the actual parameters realized by the experiment via *in situ* procedures, as discussed in Chapter 6.

2.3.3 Temperature errors

Optical tweezers trap atoms in a nearly harmonic potential [123]. As such, to a good approximation we can describe the movement of the atoms in terms of the motional Fock basis in order to capture their position and velocities. However, the state of the atoms is not pure; finite thermal excitations make the atomic motional state mixed, which can be deleterious for various reasons. Significant research over the past twenty years has been devoted to cooling atoms down to their motional ground state [124–130], and in our experiment with a little effort we are able to reach motional ground state fractions of $p_0 \approx 70\%$ purely through conventional means [85]. This fraction is then related to the approximate temperature via $p_0(T) = 1 - \exp(-(\hbar\omega)/(k_B T))$, where k_B is Boltzmann’s constant and ω is the angular trapping frequency¹. We recommend consulting Refs. [76, 117] for more details on aspects of atomic cooling that we do not discuss here.

In Chapter 3.5 we will explore a novel measurement-based cooling scheme which allows us to reach ground state fractions as high as 99%, but it is still worth it to carefully discuss the pernicious effects of temperature. Importantly, as opposed to laser errors, temperature is an inhomogeneous error source, and as such its effects are uncorrelated across the array. We largely break these effects up into three subcategories: *Doppler shifts*, *beam sampling effects*, and *interaction variations*. These effects stem from the fact that thermal atoms have finite and Gaussian position and velocity distributions, characterized respectively by standard deviations

$$\sigma_x = \sqrt{\frac{\hbar}{2m\omega} \coth\left(\frac{\hbar\omega}{2k_B T}\right)} \quad (2.8)$$

$$\sigma_v = \sqrt{\frac{\hbar\omega}{2m} \coth\left(\frac{\hbar\omega}{2k_B T}\right)}, \quad (2.9)$$

where m is the atomic mass.

¹See Ref. [76] for a nice table showing conversions between different thermal quantities.

Doppler shifts

The first, and simplest error source arises from Doppler shifts, denoted as $\Delta_{\text{dop}}^{(i)}(T)$ arising from the atoms finite thermal velocity distribution. Numerically, one samples a given velocity, \vec{v} , from Eq. 2.9, and then calculates the Doppler shift via $\delta = \vec{k} \cdot \vec{v}$, where \vec{k} is the addressing laser's wavevector. See Ref. [117] for further details of this calculation. In our system we find Doppler shifts to be a few tens of kHz at typical temperatures. We note that this description of Doppler shifts is valid in the sideband non-resolved regime, whereas in the sideband-resolved regime one accounts for Doppler shifts via the Fock-state-dependent Rabi frequency [76, 124].

Ostensibly, the velocity distribution which induces Doppler shifts could also lead to time-dependent atomic motion, affecting result we will discuss below. However, at typical microkelvin temperatures the characteristic velocity is $\sigma_v \approx 10 \text{ nm}/\mu\text{s}$, meaning roughly 100 nanometer may be traversed over the course of a $10 \mu\text{s}$ experiment. Because of this relatively small value (the context of which shall become apparent shortly), we generally consider a frozen-core approximation, and sample atomic velocities purely to calculate Doppler shifts.

Beam sampling

One of the most pernicious errors, and most often overlooked, affecting atom arrays arises from the effect of the atom sampling different laser intensities, which we call *beam sampling* $\Omega_{\text{bs}}^{(i)}(T)$. Atoms at finite temperature have a wavefunction spread, and thus on an approximately shot-to-shot basis they will sample different regions of the addressing laser. If the addressing laser has a very large waist¹, w_0 , and is well aligned, this results in a somewhat negligible effect. However, if the addressing beam is misaligned, then the atom will sample a linear slope of the Gaussian beam, and as such will be dramatically more sensitive.

The relative intensity variation for a given displacement δ_x is given by

$$\delta_I(\delta_x) = 1 - e^{-2(\delta_x/w_0)^2}. \quad (2.10)$$

Then, assuming the laser is well-aligned to the mean atomic position, we can convert from the positional Gaussian probability distribution (characterized by σ_x) to a probability distribution of the intensity².

$$P(\delta_I) = \frac{1}{(1 - \delta_I)\sqrt{-\pi\alpha \log(1 - \delta_I)}} (1 - \delta_I)^{\frac{1}{\alpha}} \quad (2.11)$$

¹Unless otherwise noted, we will always mean waist-radius when simply referring to a laser waist.

²While basic, it bears repeating that the probability density function (PDF) of a function of a random variable is *not* the same as the function applied to the PDF of the random variable. That statement is true on the level of cumulative distribution functions, so accurate PDF transformations simply require judicious use of the chain rule.

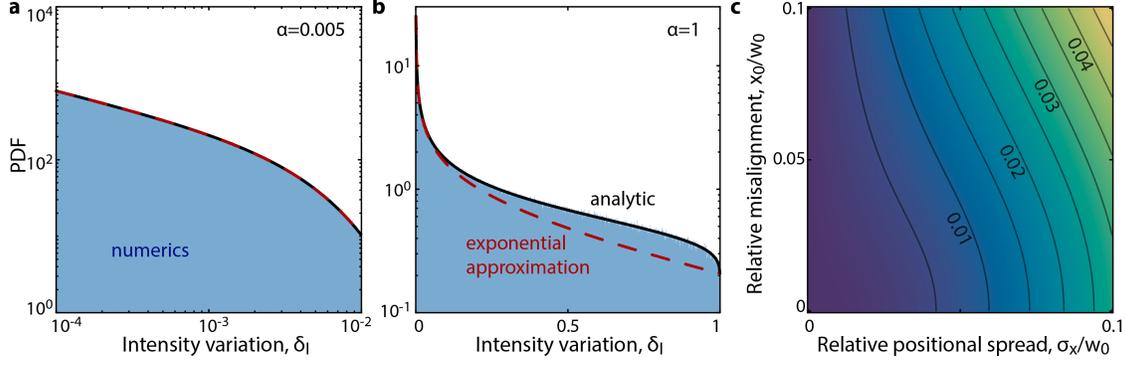


Figure 2.7: **Intensity variation from beam sampling.** **a,b**, The thermal positional spread of the atoms leads to intensity variation. Numerical simulations, assuming a Gaussian beam and thermal distribution of atoms, are in good agreement with our analytical formula (Eq. (2.11)) both when the position fluctuations are **(a)** small or **(b)** large compared to the beam waist. **c**, Beam misalignment can amplify the deleterious effects of beam sampling by leading to a larger effective standard deviation of intensity fluctuations.

$$\approx \frac{1}{\sqrt{\pi\alpha\delta_I}} e^{-\delta_I/\alpha}, \quad (2.12)$$

where $\alpha = 4\sigma_x^2/w_0^2$, and where in the second-line we have approximated $\delta_I \ll 1$ during the derivation (Fig. 2.7). For single atom Rabi oscillations with Rabi frequency Ω , this variation leads to approximate power-law decay envelope of $(1 + (t/T)^2)^{-1/4} \sim t^{-1/2}$, with $T = w_0^2/(2\Omega\sigma_x^2)$. Incorporating a finite misalignment, x_0 , the standard deviation of $P(\delta_I)$ distribution is approximately

$$\sigma_I = \sqrt{8\left(\frac{\sigma_x}{w_0}\right)^4 + 16\left(\frac{x_0}{w_0}\right)^2\left(\frac{\sigma_x}{w_0}\right)^2}. \quad (2.13)$$

Thus, the effect of this noise channel increases dramatically both as the beam waist decreases, and as the beam alignment worsens (Fig. 2.7). In our case, characteristic length scales are $\sigma_x \approx 200$ nm (along the weakly-confining axial direction), and $w_0 \approx 20$ μm . Thus, keeping the laser and atoms well-aligned (see Appendix A) is of paramount importance, as for instance a 1 μm beam misalignment can lead to nearly a 50% loss of single-atom Rabi coherence, with similarly drastic effects on the many-body fidelity. We also note that the susceptibility to this effect can be mollified through the use of super-Gaussian addressing beams [131].

Interaction variation

The position variations of the atoms do not only affect their interaction with the addressing beam, but also with each other due to the long-ranged Rydberg interaction. As such, each

atom will have its position displaced by some amount $\delta^{(i)}(T)$ which needs to be taken into account when calculating the atom-atom interaction strength. The impact of such effects falls off dramatically with distance between atoms and is further suppressed at short distances due to the Rydberg blockade mechanism, because small changes in distance do not substantively change the blockaded condition. Thus, it is most important to be aware of this effect when atoms are spaced just barely at or beyond the blockade radius from each other. This is particularly possible in two-dimensional arrays like square arrays as one must be careful of next-nearest-neighbor interactions (Fig. 11.4). Still for the one-dimensional arrays we concern ourselves with here, this effect can still lead to non-trivial fidelity loss. This effect necessarily introduces a force on the atoms, though a back-of-the-envelope calculation suggests for our typical parameters it only results in a velocity of ≈ 10 nm/ μ s, and as such we typically employ a frozen-core approximation.

2.3.4 State preparation and measurement (SPAM) errors

The survival and detection fidelities of atoms in optical tweezers are in general very high, with record values reaching or surpassing the 0.9999 level [26, 73, 85, 132]. Still, even such high values can have a dramatic effect on conclusions drawn from experimental data. We direct the reader to Ref. [34] which discusses many of these ideas for the case of one and two qubits.

Preparation errors

Typically all quantum circuits and simulations assume an input fiduciary state in the Hilbert space. When qubits are defined along the clock-Rydberg transition, this input state is for instance the *vacuum* state, or the all-zero state $|0\rangle^{\otimes N}$ of N atoms. However, our ability to prepare this state is limited by three effects: atomic survival, imperfect transfer to the clock state, and decay out of the clock state.

The atomic imaging survival in our system generally approaches the 0.9999 level [26, 85], which is nearly vacuum-limited. Thus, we expect that the probability for atom loss is relatively uniform across all times, without much regard for the dynamics affecting the atoms, such as imaging, cooling, etc. For most operations we perform, atoms are first rearranged, then imaged again to verify rearrangement has succeeded, and then there is a wait time of ≈ 200 ms before actual quantum operations. This wait time is largely limited by the time needed to change magnetic fields. However, after we perform the quantum operation, there is < 1 ms of wait time before the final readout image; because of this separation of timescales, we can assume that any atom loss occurs *before* the quantum

operation. Such losses leave unidentified holes in the array, which can have dramatic ramifications for quantum dynamics; we shall explore this effect in great detail in Chapter 3 where we detect these errors through *erasure conversion*.

Beyond atomic survival, preparation is also limited by imperfect clock state preparation. We generally are able to drive the clock transition coherently with ≈ 0.998 fidelity [28, 39], limited by frequency noise on the clock laser and atomic temperature. Further, we can supplement clock state preparation with incoherent pumping [34], which enables close to survival-limited preparation (though atoms which are pumped are also likely prepared in a much hotter thermal state, which we do not explore the implications of).

Once prepared, there is generally a wait time of a few milliseconds before Rydberg excitation to account for shutter opening times etc. and so some atoms will decay out of the metastable clock state in that time. For bosonic strontium the lifetime of the clock state is conjectured to be thousands of years [107], but in practice it is limited by Raman scattering from the high-intensity trapping light [133]. Specifically, the tweezer light at 813 nm can induce a far off-resonant coupling to the 3S_1 excited state, which can then rapidly decay to any state in the 3P_j manifold. The rate of scattering is roughly proportional to the trap depth (i.e. tweezer intensity), which we find [33] to be $\approx 250(30) \text{ s}/\mu\text{K}$. At our normal trap depths this amounts to a lifetime on the order of $\approx 1 \text{ s}$.

Thus, we generally find an overall preparation fidelity of ≈ 0.997 per atom. Preparation failure is a leakage error, for which we make the atom inert. There are two possibilities, that the preparation error was from atom loss (which would be dark to our final imaging), or was from imperfect clock state transfer (which would likely be bright to our final imaging); in the former case, we switch the inert ‘qubit’ state to a $|1\rangle$ to have it be properly accounted for during readout.

Measurement errors

When the quantum state is measured, it collapses into a distinct basis state with a probability given by the state vector. However, at this point a classical error channel can distort the resulting bitstring. We distinguish two types of measurement errors: Rydberg detection (autoionization) errors and imaging errors.

Autoionization errors In order to distinguish Rydberg atoms from clock atoms, we apply a short and powerful autoionizing pulse [26, 34]. This pulse ionizes Rydberg atoms with a fidelity in excess of $F_A = 0.999$, though we note care must be taken to minimize wait times between the end of Rydberg excitation and the start of the autoionization pulse otherwise

Rydberg decay can mitigate the autoionization fidelity. If the autoionization pulse fails, then it is essentially a classical amplitude-damping channel mapping $|1\rangle \rightarrow |0\rangle$.

Imaging errors The state of the atoms is read out through fluorescence imaging by ideally scattering many photons quickly, collecting them on a camera, and weighting them with a precalibrated point-spread-function (PSF) mask¹. The basics of this process are discussed in detail in Ref. [34], but essentially it amounts to two different detection fidelities: the false positive fidelity (F_0) for an atom in state $|0\rangle$ to be readout as a ‘0’, and the false negative fidelity (F_1) for an atom in state $|1\rangle$ to be readout as a ‘1’. These fidelities are typically balanced by using a threshold based approach [73, 99], accounting for the tweezer filling fraction [34, 73], though we note more advanced imaging methods involving neural networks may improve imaging fidelity [134, 135].

To account for both imaging and autoionization errors, we employ a transfer function based approach. We first take the absolute square of the quantum state vector to form the population vector. Then, for each individual qubit in the array (see Appendix B for details), we apply the transfer matrix

$$\mathbf{M} = \begin{bmatrix} F_1 & (1 - F_A)F_1 + F_A(1 - F_0) \\ 1 - F_1 & F_AF_0 + (1 - F_A)(1 - F_0) \end{bmatrix}. \quad (2.14)$$

Note that if an atom is marked as inert, then $F_A \equiv 1$. The resulting population vector then is an accurate model for experimental observables.

2.3.5 Spontaneous decay errors

The Rydberg state is inherently short-lived, with lifetimes on the order of 10s of μs , where the lifetime increases with increasing principal quantum number [76]. Due to its highly excited nature, the Rydberg state has substantial branching ratio to decay along a number of different pathways. As a simplification, we generally consider two possible channels: 1) spontaneous decay down to a low-lying energy manifold which is detected in our final image (bright-state decay), 2) blackbody radiation induced decay to a manifold which is dark to our final image, such as nearby Rydberg levels (dark-state decay). These processes

¹In our work, besides the specific case of erasure conversion in Chapter 3 we always perform all images with the same imaging parameters, i.e. the image used prior to rearrangement is the same as the image used for final readout. However, this is not really the most ideal strategy, as different images have different requirements and sensitivities to detection fidelities and atomic survivals. For instance, in the final readout image we could stand to gain by ignoring atom survival and only maximizing photons scattered to improve detection.

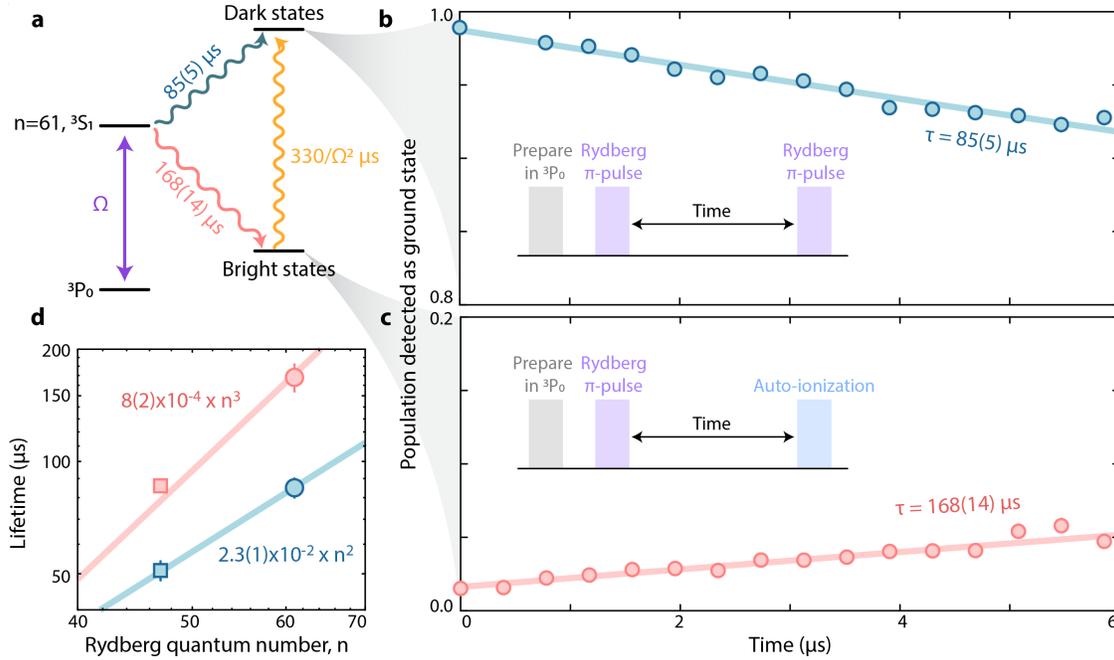


Figure 2.8: Rydberg lifetime. **a**, Pseudo-level diagram of the Rydberg dynamics and associated decay channels. During Rydberg evolution with Rabi frequency Ω , the Rydberg atom can decay into either a set of states which is ‘bright’ to the imaging process (including both the erasure images and the final detection images), e.g. states like $5s5p\ ^3P_2$, or into states which are ‘dark’ to the imaging process, e.g. nearby Rydberg states. A small percentage of decays into ‘bright’ states can go directly into 3P_0 where they can be re-excited by the Rydberg driving; note that such decays are dark in the erasure image, but bright in the final detection image. **b**, Measurement of the dark state decay lifetime, measured by performing a π -pulse on the Rydberg transition, waiting a variable amount of time, and then returning atoms to the ground state (inset). **c**, Measurement of the bright state decay lifetime, measured by performing a Rydberg π -pulse, waiting, and then performing an auto-ionization pulse to destroy any remaining Rydberg or dark state excitations. Markers are experimental data where error bars are often smaller than the marker sizes, and solid lines represent exponential fits. **d**, Aggregating data from our own study (Ref. [26], circles) with that of Ref. [136] (squares) shows good agreement to the expected scaling [76] of the Rydberg decay lifetimes with the principal quantum number. Note that bright state decay and dark state decay have different expected scaling coefficients with Rydberg quantum number, n , of n^3 and n^2 , respectively [76].

are broadly summarized in Fig. 2.8. We note that all numerical values below are measured for a principal quantum number of $n = 61$. Different decay mechanisms have different scaling with n (see Table 4.5 in Ref. [76] for a summary), but we note all our findings are consistent with that of Ref. [136] (Fig. 2.8d).

We simulate spontaneous decay processes using the quantum jump method [114] which

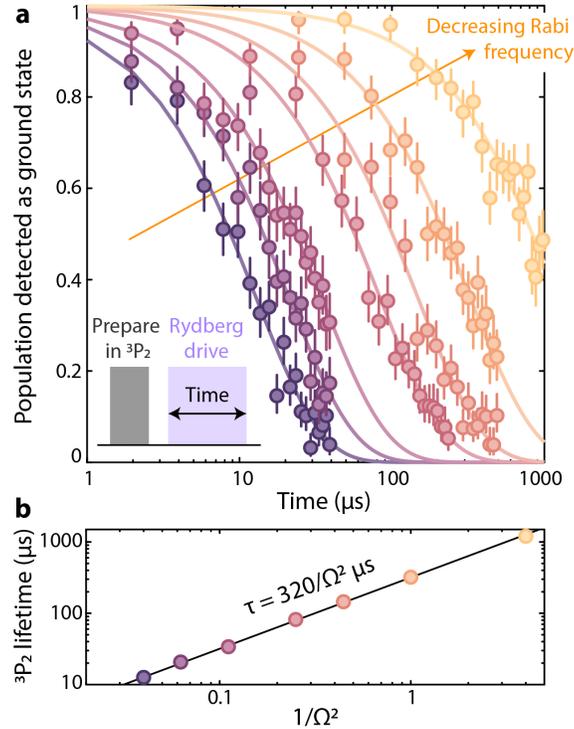


Figure 2.9: 3P_2 **photo-ionization**. **a**, We prepare atoms into 3P_2 (a bright state), and then perform continuous Rydberg driving. Atoms are lost from the trap at a rate which increases with increasing Rabi frequency. **b**, The lifetime of atoms in 3P_2 scales inversely with the square of the Rabi frequency (i.e. scales inversely with the intensity of the Rydberg beam). We attribute this to a photo-ionization process which can convert bright state decay into dark state decay through prolonged Rydberg excitation, as shown in **a**.

employs non-Hermitian evolution in concert with probabilistic decay events. For each time interval dt , the decay channel $|i\rangle \rightarrow |f\rangle$ occurs with probability $\Gamma_{if}|\langle i|i\rangle\psi(t)|^2dt$, where $|\psi(t)\rangle$ is the quantum state at time t . If a quantum jump happens, we will apply the projection operator $|f\rangle\langle f|$ to the state.

We find the following lifetimes for the various decay channels (Fig. 2.8: 1) dark state decay - $\tau = 85(5)\mu s$; 2) decay to the $5s5p$ 3P_J manifold - $\tau = 168(14)\mu s$ and branching ratio 10%, 30%, 60% for $J = 0, 1, 2$. Further, atoms in 3P_1 and 3P_2 are photoionized by the Rydberg drive (Fig. 2.9), with a Rabi-frequency-dependent timescale of $\tau = 320\mu s/(\Omega/\text{MHz})^2$.

We note that we do not account for a potential major error source, namely correlated spontaneous decay, as has been observed in Rydberg-dressed lattices [137], and which we do see some signatures of. Essentially, when a Rydberg atom via blackbody radiation to a dark state, the decay wavelength is much longer than the array lattice spacing. This leads to an avalanche phenomena where *all* atoms simultaneously decay. This effect is pronounced

when the lattice spacing is tighter. We do not explicitly take into account this effect in our error model, but note that doing so is likely not too difficult. To do so, if we detect a quantum jump has occurred during evolution we would perform conditional nearest-neighbor checks. Specifically, we would introduce a new random number parameter, such that if there is a dark state decay at site i , we draw two random numbers for sites $i - 1$ and $i + 1$. If either of these random numbers is lower than a constant, q , then we project these sites into dark states. If the projection occurs, we repeat the same process for the $i - 2$ and $i + 2$ sites (and so on) to capture the proliferation of dark states, and so on. In this scheme q would likely be calibrated directly against *in situ* many-body measurements.

Erasure conversion

Having described many of the error sources facing Rydberg atom array quantum simulators, we now turn to discuss how to eliminate many of them using a new tool in the quantum toolbox: *erasure conversion*. The definition of erasure conversion is still somewhat in flux, but fundamentally it means the direct detection of errors as they occur *in situ* during a quantum algorithm, so that they may be heralded or corrected for. The specification that this detection is *direct* is key: we are not detecting errors via syndrome measurements, as is standard in quantum error correction, but instead by reading off a per-qubit signal which flags if that qubit has suffered an error. Our demonstration of erasure conversion is the first of its kind (alongside Ref. [75]) following the theoretical proposal of Ref. [68] (and early discussions in Ref. [67]), and could have significant implications both for near term quantum applications and for reducing the pseudo-threshold needed to achieve error correction [68].

In this chapter, we will first describe the importance of erasure conversion, and how it is accomplished in atom arrays. We then describe several case studies of its application, such as improving the preparation fidelity of two-qubit Bell states, and for improving quantum simulation fidelities for large-range entangled states [26]. Using the principles of erasure conversion, we then demonstrate a erasure-based method for cooling atoms to their motional ground state [32], and show a measurement- and feedback-based method for improving atom array loading, called dark-state enhanced loading [33].

3.1 LOCATING ERRORS

Errors are rampant in quantum systems, and Rydberg atom arrays are no different. A common, and lofty, goal of the field is to achieve so-called *error correction*, where many physical qubits are aggregated in order to perform the functions of a single *logical* qubit at a much reduced error rate. While quantum error correction schemes can be quite involved, their simplest classical analogs are quite intuitive. For instance, the simplest classical error correcting code is the repetition code.

Say there is a single bit of information which is subject to a noisy channel. To improve the transmission fidelity we can simply copy the bit several times, producing, say, 9 identical bits total, and then send all the resultant bits through the channel. The receiver on the other end then evaluates all simultaneously and uses majority vote to determine the bit value (i.e. if eight of the transmitted bits are a “1”, and only one is a “0”, then it is more likely than not that the true value is a “1”). This approach, while relatively memory inefficient, is quite effective classically, and spiritually underlies many of the decoding mechanisms for quantum error correction techniques.

Still, let’s imagine error rates are slightly higher, and instead when we send our 9 identical bits through the noisy channel the result is that five are “1” and four are “0.” We could still apply majority vote, but the statistical significance is greatly reduced. For this reason, we ideally want to have a method of pruning some of the “bad” bits so that we are less affected by their errors. If we can use some auxiliary information to be sure that all the “0” instances came from erroneous transmissions, then we regain high statistical certainty in the aggregate (logical) value. This process of locating where errors have occurred is precisely the point of erasure conversion. Theoretical proposals for erasure conversion [68] discussed it explicitly in the context of improving the ability of near term quantum computers to implement error correction techniques.

It is simplest to first consider the most classical version of erasure correction, namely atomic rearrangement. When atoms are loaded into an optical tweezer array they do so stochastically, only filling a given site in the array with roughly a 50% probability. Consider running an algorithm on N atoms, without implementing rearrangement protocols. In such a case, the probability of success will be at best 2^{-N} , as any missing atom guarantees failure. Seemingly daunting, we know this impedance is immaterial, given that it is straightforward to image the array before performing the operation, and rearrange [81, 89] to ensure a defect free arrangement of our N atoms. In other words, we have prepared an array of atoms, then it has undergone a noisy channel of stochastic filling, then we directly measure the “errors,”

by imaging where the atoms are located, and subsequently we correct via rearrangement.

To be clear, there are distinctions and caveats to be drawn from this simple example, which in some sense is an example of classical erasure conversion. We employ the following definitions: *Classical erasure conversion* - all qubits (bits) are readout (projected) for errors; *Quantum erasure conversion* - only erroneous qubits are readout, while non-erroneous qubits are not projected; *Mid-circuit erasure conversion* - the readout (and potentially correction) is fast compared to the qubit coherence time. Clearly, atomic rearrangement falls into the category of ‘classical’ erasure conversion, per our definitions. Still, it helps to convey the basic point: if we can image where errors take place, and can do so in such a way so as not to disturb error-free qubits we can potentially dramatically improve device performance (and hasten the arrival of error corrected quantum systems [68]).

3.2 AN ATOMIC ERASURE CONVERSION IMPLEMENTATION

Here we describe how we perform the erasure imaging which allows us to detect site-resolved leakage errors out of the qubit subspace spanned by the clock state and the Rydberg state.

Our basic protocol relies on the long-lived nature of the clock state, and the fact that both improper preparation into the clock state or bright state decay (see Section 2.3.5) out of the qubit manifold will typically result in erroneous atoms cascading into 1S_0 . To exploit this, we develop a very fast atom detection technique (previously showcased for freely propagating lithium atoms [138]) which allows us to detect atoms in 1S_0 in a $\approx 20 \mu\text{s}$ timescale (Fig. 3.1a). The basic imaging technique, which is covered in great detail in Ref. [138], essentially relies on scattering a large number of photons very quickly, and then overcoming intrinsic camera noise through a combination of low-pass and binary filters on the camera signal.

We highlight the characteristic features of this imaging scheme experimentally. We show in Fig. 3.1b the survival probability of atoms in 1S_0 as a function of imaging time. After $4 \mu\text{s}$, more than 80% of the atoms are lost. However, the number of detected photons continues to increase: even though the kinetic energy of the atoms is too large to keep them trapped, their mean position remains centered on the tweezers. In order to both avoid any extra heating coming from the imaging beams and optimize imaging fidelity, we shine two identical counter-propagating beams with crossed π -polarization and Rabi frequencies $\Omega/2\pi \cong 40 \text{ MHz}$ on the $^1S_0 \rightarrow ^1P_1$ transition. This minimizes the net force on an atom, and the crossed polarization avoids intensity interference patterns. After $\approx 24 \mu\text{s}$, the atomic spread becomes too large and the number of detected photons plateaus.

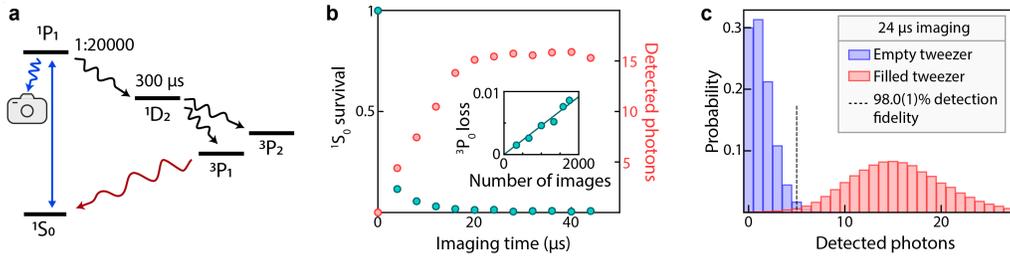


Figure 3.1: **Fast imaging on the erasure detection subspace.** **a**, Sketch of the involved energy levels. We detect atoms in 1S_0 by strongly driving the $^1S_0 \leftrightarrow ^1P_1$ transition. **b**, Survival of atoms in 1S_0 (green) and number of detected photons (red) as a function of the imaging time. We observe an increase of detected photons whereas the atoms are already lost: even though the kinetic energy of the atoms is too large to keep them trapped, their mean position remains centered on the tweezers thanks to the use of two counter-propagating beams with equal power. After $\approx 24 \mu\text{s}$, the atomic spread becomes too large to measure a significant increase in detected photons. Inset: losses from 3P_0 as a function of time, expressed in number of fast images. The survival probability of an atom in 3P_0 is 99.99954(12)% for one image, consistent with its 5 second lifetime. **c**, Typical histograms of the number of detected photons for $24 \mu\text{s}$ imaging. Using a slow, high-fidelity image prior to the fast image, we can detect if a tweezer is empty (blue) or filled (red). The typical detection fidelity which corresponds to equal error probability in detecting absence or presence of an atom is 98.0(1)%.

We note that in our implementation, we explicitly scatter a very large number of photons very quickly such that cooling the imaged atoms is not possible and they are almost inevitably lost during the imaging process¹. While unfortunate, in this first experimental demonstration it is not too much of a bother, given that the only atoms that will be lost in this way are those for which errors had occurred; specifically these errors are leakage errors out of the qubit subspace, which are anyway difficult to correct without replacing the atom. In the future, it will be necessary to either perform these ultrafast images non-destructively for instance using cavity-assisted readout [139], or to provide a steady stream of fresh atoms via continuous loading techniques [33, 84, 140, 141].

The obtained detection histogram is shown in Fig. 3.1c. We present the results both for empty (blue) and filled (red) tweezers, which we achieve by first imaging the atoms using usual, high survival imaging for initial detection in a 50% loaded array, then perform the fast image. We obtain a typical detection fidelity of 0.980(1) of true positives and true negatives, limited by the finite probability for atoms in 1P_1 to decay into 1D_2 (Fig. 3.1a).

¹This is in contrast to our standard method of imaging, described at length in Ref. [76] which continually cools atoms during imaging such that their survival is nearly vacuum-limited.

We note this fidelity is somewhat dependent on the array spacing: atoms can travel over relatively large distances during the violent imaging process, which we conjecture mixes their imaging point-spread-functions more than in standard imaging. We are able to see fidelities in excess of 0.99 by spreading atoms further apart prior to imaging.

This imaging scheme is sufficiently fast to avoid perturbing atoms in 3P_0 , as measured by losses from 3P_0 as a function of imaging time (Fig. 3.1b, inset). We fit the data (circles) using a linear function (solid line), and obtain a loss of 0.0000046(12) per image, consistent with the lifetime of the 3P_0 state [33] of ≈ 5 s for the trap depth of 45 μK used during fast imaging.

We note that current limitations to this the scheme are not yet fundamental, and there are a number of technical improvements which could be made. First, the camera we use (Andor iXon Ultra 888) has a quantum efficiency of $\approx 80\%$, which has been improved in some recent models. Further, we currently collect photons only from one direction, when in principle photons could be collected from both objectives [38]. This would improve our estimated total collection efficiency of $\approx 4\%$ by a factor of 2, leading to faster imaging times with higher fidelity (as more photons could be collected before that atoms were ejected from the trap). Finally, but not exhaustively, the fidelity may be substantially improved by actively repumping the 1D_2 state back into the imaging manifold so as to not effectively lose any atoms via this pathway [142].

3.3 CASE STUDY: IMPROVING BELL STATE GENERATION

Having discussed the physical implementation of our erasure conversion scheme, we now demonstrate it is practically useful on an important problem of interest, namely Bell state generation resulting in fidelities competitive with other state-of-the-art platforms [143–146]. Bell states are the smallest example of entangled states, with applications in computing [15], metrology [147], and communication [148].

First, we reiterate some of the basic physics at play. To controllably generate entanglement between atoms, we employ Rydberg interactions [94, 149, 150]. When two atoms in close proximity are simultaneously excited to high-lying electronic energy levels, called Rydberg states, they experience a distance-dependent van der Waals interaction $V = C_6/r^6$, where r is the interatomic spacing, and C_6 is an interaction coefficient. If the Rabi frequency, Ω , which couples the ground, $|g\rangle$, and Rydberg, $|r\rangle$, states is much smaller than the interaction shift, $\Omega/V \ll 1$, the two atoms cannot be simultaneously excited to the Rydberg state (Fig. 3.2b, right), a phenomena known as Rydberg blockade. In this regime, the laser

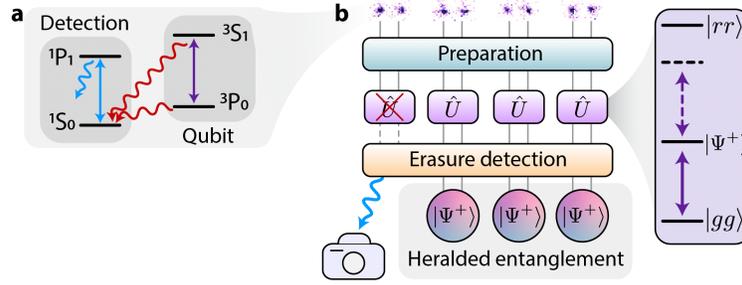


Figure 3.2: **Erasure conversion for high-fidelity entanglement.** **a**, Level structure used in this work. We distinguish two subspaces: a qubit subspace in which the atoms interact via their Rydberg states, and a measurement subspace used to detect leakage errors from the qubit subspace with single-site resolution, realizing erasure conversion. **b**, Sketch of the erasure conversion scheme, as applied to Bell pair generation. After arranging atoms into pairs (top) we prepare them in $|g\rangle$, and entangle them via the Rydberg blockade mechanism (right), denoted by a unitary operation $\hat{U}(t)$. Immediately afterwards, we auto-ionize atoms in $|r\rangle$, effectively projecting the populations of the Bell states, and follow with a fast erasure conversion image to detect leakage out of the qubit subspace during the preparation or evolution periods. This is followed by the final detection of atoms in $|g\rangle$, yielding two separate, independent images. We can discard data from pairs where atoms are detected in the erasure-error image, termed erasure-excision in the following.

drives a unitary operation, $\hat{U}(t)$, that naturally results in the two atoms forming a Bell state, $|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|gr\rangle + |rg\rangle)$, between the ground and Rydberg states (Fig. 3.2b).

3.3.1 Protocol

This Bell state generation has several major practical limitations. Of particular interest here are leakage errors to the absolute ground state, 1S_0 . The first error of this type is imperfect preparation of atoms in $|g\rangle$ prior to applying $\hat{U}(t)$. The second arises from decay out of the Rydberg state along multiple channels. We distinguish decay into ‘bright’ states, which we can image, and ‘dark’ states, which are undetected (Fig. 2.8). The former primarily refers to low-lying energy states which are repumped to 1S_0 as part of the imaging process or decay to 1S_0 via intermediate states, while the latter mainly consists of nearby Rydberg states accessed via blackbody radiation.

Our general procedure is shown in Fig. 3.2 (further detailed in Fig. 3.3). We first rearrange [81, 89] atoms into pairs, coherently transfer them to $|g\rangle$, and then perform the entangling \hat{U} operation. Immediately after, we auto-ionize the atoms to project the populations of the resultant state.

We then perform the fast erasure image; any atoms which are detected are concluded to be the result of some leakage error process. Thus, we obtain two separate images characterizing

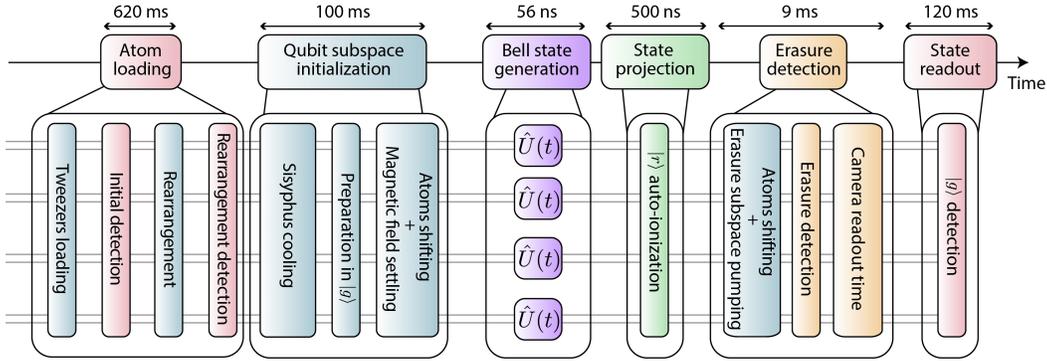


Figure 3.3: **Experimental sequence.** Sketch of the experimental sequence including the erasure detection for the Bell state generation experiment. We start by loading the atoms into the desired geometry, then initialize the atoms in $|g\rangle$, perform the Bell state generation or quantum simulation, and finally read out by auto-ionizing atoms in $|r\rangle$ and imaging atoms in $|g\rangle$. We utilize a single erasure detection, placed after auto-ionizing atoms in $|r\rangle$.

a single experimental repetition, the final image showing the ostensible result of \hat{U} , and the erasure image revealing leakage errors with single-site resolution.

We note that this work is not a form of mid-circuit detection as no superposition states of $|g\rangle$ and $|r\rangle$ exist at the time of the erasure image. Instead, our approach is a noise mitigation strategy via erasure-excision, where experimental realizations are discarded if erasures are detected. Importantly, in contrast to other leakage mitigation schemes previously demonstrated in matter-based qubit platforms [151–153], we directly spatially resolve leakage errors in a way which is decoupled from the performed experiment, is not post-selected on the final qubit readout, and does not require any extra qubits to execute.

However, the coherence between $|g\rangle$ and $|r\rangle$ can in principle be preserved during erasure detection for future applications; in particular, we see no significant difference in Bell state lifetime with and without the imaging light for erasure detection on (Fig. 3.5). We also expect long-lived qubits encoded in $|g\rangle$ to be unperturbed by our implementation of erasure conversion [32, 68, 75, 104, 113], as we for instance demonstrate in Section 3.5 for the case of a long-lived qubit encoded in the motional manifold.

3.3.2 Results

With a procedure for performing erasure conversion in hand, we now describe its impact on Bell state generation. Experimentally, we only obtain a lower-bound for the Bell state generation fidelity [34] (Fig. 3.6); the difference of this lower bound to the true fidelity is

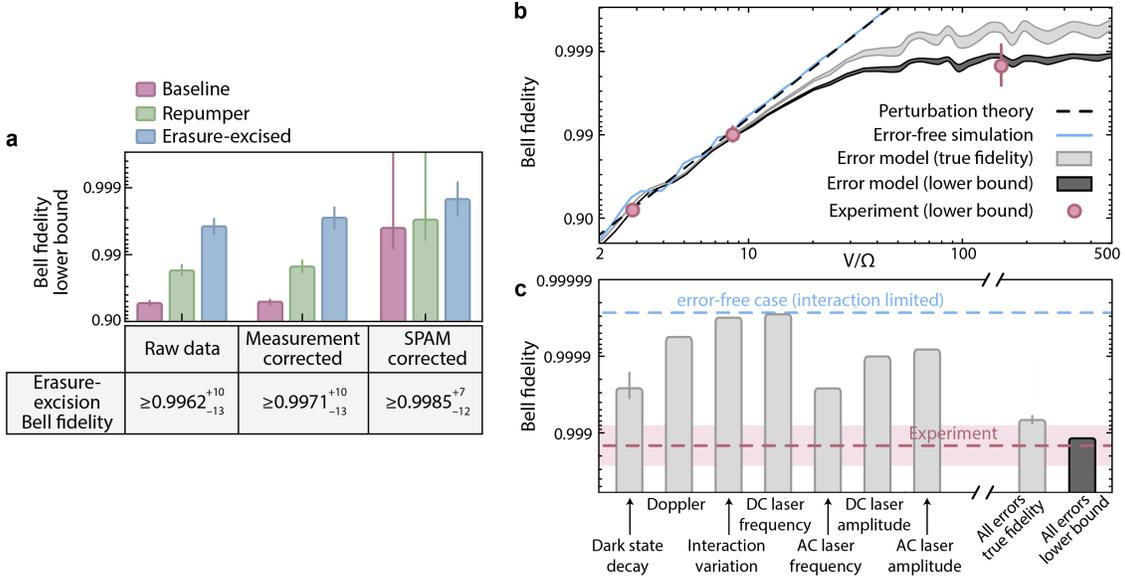


Figure 3.4: **Predicting infidelities at the 10^{-3} level.** **a**, Lower bounds for Bell state fidelities with (blue) and without (red) the erasure-excision, and using incoherent pumping to reduce preparation errors instead of erasure-excision (green). We present the results for the raw data, corrected for measurement errors, and corrected for state preparation and measurement errors (SPAM). All data is averaged over eight pairs of atoms which are excited in parallel. Error bars represent a 68% confidence interval. **b**, SPAM-corrected Bell state fidelity as a function of the ratio of interaction energy and Rabi frequency, V/Ω . Error-free simulations (blue line) show fidelities continually increase with increasing V/Ω , in agreement with results from perturbation theory (dashed line). For large enough interaction strength ($V/\Omega > 50$), other error sources become dominant, and we employ a noisy open system dynamics simulation from which we obtain an estimate of the true fidelity (light grey fill) and for the lower bounding procedure used in experiment (dark grey fill). We find good agreement between simulation and experimental results with erasure-excision (red markers). **c**, Predicted Bell state fidelity for $V/\Omega = 140$ from simulations turning on a single noise term at a time. Dominant limitations come from laser frequency and intensity noise, as well as decay of the Rydberg state into dark states. We also show the results when taking into account all errors, for both the true fidelity and the lower bound estimation (right). The lower-bound significantly underestimates the true fidelity. Shaded areas in **b** and error bars in **c** represent the standard deviation of the mean over 5000 trajectories.

discussed further below.

We first coherently transfer atoms to $|g\rangle$ as described before, and then consider three scenarios (Fig. 3.4a). In the first, as a baseline we perform the entangling unitary \hat{U} without considering any erasure detection results (red bars). In the second, we excise data from any pairs of atoms with an observed erasure error (blue bars). Finally, we compare against another strategy for mitigating preparation errors through incoherent repumping [34], but without erasure detection (green bars). Notably, the raw value for the Bell state lower-

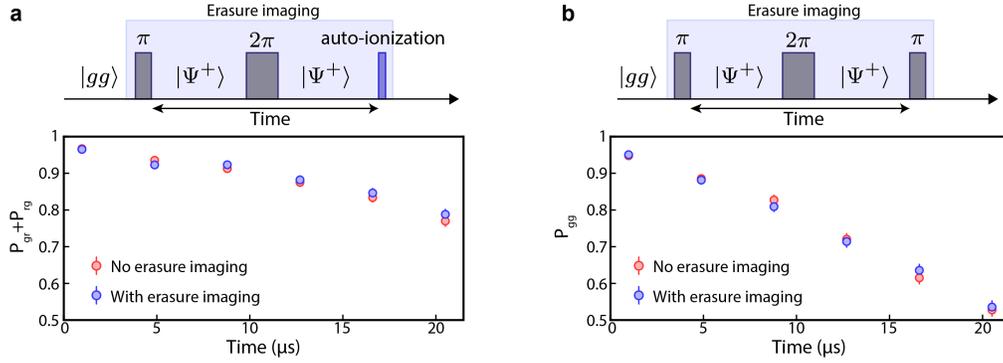


Figure 3.5: **Coherence preservation during erasure detection.** **a,b**, We prepare the Bell state $|\Psi^+\rangle$, and measure the relevant populations for Bell fidelity extraction $P_{gr} + P_{rg}$ and P_{gg} (after an extra π pulse) as a function of holding time. We perform a 2π pulse in the middle of the holding time to get rid of dephasing effect due to e.g. Doppler effect. We present the results with (blue) and without (red) performing the erasure imaging during the holding time. We observe no significant difference between the two conditions, which suggests that the erasure detection imaging light, in principle, does not destroy the coherence of the Bell state.

bound with erasure-excision is $\geq 0.9962(+10/-13)$, significantly higher than with the other methods. This difference mainly comes from erasure excision of preparation errors and, to a much lower degree, Rydberg decay. These contribute at the level of $\sim 5 \times 10^{-2}$ and $1.2(3) \times 10^{-4}$, respectively.

Correcting for final measurement errors, we find a lower bound of $\geq 0.9971(+10/-13)$, which quantifies our ability to generate Bell pairs conditioned on finding no erasure events. To quantify the quality of the Rydberg entangling operation $\hat{U}(t)$ itself, we further correct for remaining preparation errors that are not detected in the erasure image, and find a state preparation and measurement (SPAM) corrected lower bound of $\geq 0.9985(+7/-12)$.

To our knowledge, these bare, measurement-corrected, and SPAM-corrected values are, respectively, the highest two-qubit entanglement fidelities measured for neutral atoms to date, independent of the means of entanglement generation. While Bell state generation as demonstrated here is not a computational two-qubit quantum gate, our results are indicative of the fidelities achievable in Rydberg based gate operations.

3.3.3 Error modelling

Importantly, we understand remaining errors in the entangling operation as well the nature of detected erasure errors from our detailed *ab-initio* error model simulation for SPAM-

corrected fidelities (Fig. 3.4). Parameters are not fine-tuned to match the measured Bell state fidelity, and the model equally well describes results from many-body quench experiments [29].

We identify limited interaction strength as a dominant effect that restricted SPAM-corrected entanglement fidelities in our previous work [34] (Fig. 3.4b); in particular, one major difference here is that we operate at smaller distance and hence larger V/Ω . In line with experimental data (red markers), fidelities at large distances are limited to $F_{\text{Bell}} \leq 1 - \frac{5}{8}(\Omega/V)^2$ obtained from perturbation theory (black dashed line).

For strong enough interaction, $V/\Omega > 50$, corresponding to distances $r < 3 \mu\text{m}$, other error sources become limiting. In this short-distance regime, the experimental SPAM-corrected fidelity lower-bound is in good agreement with the error model prediction of $\geq 0.99881(3)$ (dark grey fill).

Our error model results show that the lower bound procedure significantly underestimates the true fidelity (light grey fill), found to be $0.99931(6)$. This effect arises because the lower bound essentially evaluates the fidelity of \hat{U} by a measurement after performing \hat{U} twice, meaning particular errors can be exaggerated. Given the good match of the error model and experimental fidelity lower bounds, we expect this effect to be present in experiment as well, and to underestimate the true SPAM-corrected fidelity by about 5×10^{-4} .

Atoms can decay from 3P_0 between state preparation and Rydberg excitation, which is 60 ms to allow time for magnetic fields to settle. In previous work [34], we supplemented coherent preparation with incoherent pumping to 3P_0 immediately prior to Rydberg operations. However, during the repumping process, atoms can be lost due to repeated recoil events at low trap depth, which is not detected by the erasure image, and thus can lower the bare fidelity. Even with SPAM-correction of this effect, we expect the fidelity with repumping to be slightly inferior due to an increased atomic temperature for pumped atoms.

As to the nature of the detected erasure errors for the Bell state generation, we find that preparation errors contribute the vast majority of erasure events compared to bright Rydberg decay, and excising them has a more significant impact on reducing infidelities. In particular, application of \hat{U} lasts for only ~ 59 ns, which is significantly shorter than the independently measured bright state decay lifetime of $168(14) \mu\text{s}$ (Fig. 2.8). The error model described in Fig. 3.4b suggests that excising such errors results in an infidelity reduction of only $1.2(3) \times 10^{-4}$. Conversely, preparation errors account for $\sim 5 \times 10^{-2}$ infidelity per pair due to the long time between preparation in $|g\rangle$ and Rydberg excitation (Fig. 3.3). Hence, the gains in fidelity from erasure-conversion mainly come from nearly eliminating all

preparation errors, which has the added benefit of significantly reducing error bars on the SPAM-corrected values. Still, SPAM-corrected values might also benefit from the small gain in eliminating the effect of bright state decay, and from avoiding potential deleterious effects arising from higher atomic temperature in the repumper case.

The remaining infidelity is a combination of multiple errors. In Fig. 3.4c, we report an error budget for the most relevant noise source contributions to the Bell state infidelity at the experimentally chosen $V/\Omega = 140$. Frequency and intensity laser noise are dominant limitations, but could be alleviated by improving the stability of laser power, and reducing its linewidth, for instance via cavity filtering [154]. Eliminating laser noise completely would lead to fidelities of ~ 0.9997 in our model. The other major limit is Rydberg state decay into dark states, which cannot be converted into an erasure detection with our scheme. This decay is mostly blackbody induced [68, 155], and thus could be greatly reduced by working in a cryogenic environment [86], leaving mostly spontaneous decay that is bright to our erasure detection. Accounting for these improvements, it is realistic that Rydberg-based Bell state generation in optical tweezers arrays could reach > 0.9999 fidelity in the coming years.

3.3.4 Extraction of the Bell state fidelity

In order to extract the Bell state fidelities quoted in the main text, we use a lower bound method [34] which relies on measuring the populations in the four possible states P_{gr} , P_{rg} , P_{gg} and P_{rr} during a Rabi oscillation between $|gg\rangle$ and $|\Psi^+\rangle$. The lower bound on Bell state fidelity is given by:

$$F_{\text{Bell}} \geq \frac{P_{gr+rg}^\pi}{2} + \sqrt{\frac{\sum_i (P_i^{2\pi})^2 - 1}{2} + P_{gr}^\pi P_{rg}^\pi}, \quad (3.1)$$

where $P_i^{2\pi}$ are the measured probabilities for the four states at 2π , and P_{gr+rg}^π is the probability $P_{gr} + P_{rg}$ measured at π . In order to measure these probabilities with high accuracy, we concentrate our data-taking around the π and 2π times (Fig. 3.6a), and fit the obtained values using quadratic functions $f(t) = p_0 + p_1(t - p_2)^2$. We first detail the fitting method, then how we obtain the four probabilities, and finally the extraction of the Bell state fidelity from these.

Fitting method

We perform a fit which takes into account the underlying Beta distribution of the data, and prevents systematic errors arising from assuming a Gaussian distribution of the data. The aim of the fit is to obtain the three-dimensional probability density function $Q(p_0, p_1, p_2)$

of f , using each experimental data point i defined by its probability density function $\mathcal{P}_i(x)$, where x is a probability. In order to obtain a particular value of $Q(\tilde{p}_0, \tilde{p}_1, \tilde{p}_2)$, we look at the corresponding probability density function value $\mathcal{P}_i(f(t_i))$ for each data point i , where $f(t_i) = \tilde{p}_0 + \tilde{p}_1(t_i - \tilde{p}_2)^2$, and assign the product of each $\mathcal{P}_i(f(t_i))$ to the fit likelihood function:

$$Q(\tilde{p}_0, \tilde{p}_1, \tilde{p}_2) = \prod_i \mathcal{P}_i(f(t_i)). \quad (3.2)$$

We repeat this for various $[\tilde{p}_0, \tilde{p}_1, \tilde{p}_2]$.

The result of such fitting method is shown in Fig. 3.6b (black line), where we present $f(t) = p_0 + p_1(t - p_2)^2$ for $[p_0, p_1, p_2]$ corresponding to the maximum value of $Q(p_0, p_1, p_2)$. We emphasize that this results in a lower peak value than a standard fitting procedure which assumes underlying Gaussian distributions of experimentally measured probabilities (red line). Choosing this lower peak value eventually will provide a more conservative, but more accurate value for the Bell state fidelity lower bound than the naïve Gaussian approach.

Obtaining the four probability distributions

Our method to obtain the probability density functions of the four probabilities at π and 2π times ensures both that the sum of the four probabilities always equals one, and that their mutual correlations are preserved. We first extract the Beta distribution of P_{rr} by gathering all the data around the π and 2π times (Fig. 3.6c). In particular, the mode of the obtained Beta distribution at π is $P_{rr} \simeq 0.0005$. The distribution of P_{gr+rg} and P_{gg} are obtained by fitting the data in the following way. We perform a joint fit on P_{gr+rg} using a fit function $f_1(t)$, and on P_{gg} using a fit function $f_2(t)$. The fit functions are expressed as:

$$f_1(t) = p_0 + p_1(t - p_2)^2, \quad (3.3)$$

$$f_2(t) = 1 - p_0 - P_{rr} - p_1(t - p_2)^2, \quad (3.4)$$

which ensures that the sum of the four probabilities is always equal to 1. We then calculate the joint probability density function $Q_{1,2}(p_0, p_1, p_2)$ of both f_1 and f_2 using the method described above. In particular:

$$Q_{1,2}(\tilde{p}_0, \tilde{p}_1, \tilde{p}_2) = \prod_i \mathcal{P}_i^{gr+rg}(f_1(t_i)) \prod_i \mathcal{P}_i^{gg}(f_2(t_i)), \quad (3.5)$$

where \mathcal{P}_i^{gr+rg} (\mathcal{P}_i^{gg}) is the probability density function associated with P_{gr+rg} (P_{gg}) for the i -th experimental data point. In particular, we impose that $p_0 \leq 1 - P_{rr}$ to avoid negative probabilities. We show the resulting $Q_{1,2}(p_0, p_1, p_2)$ in Fig. 3.6d as two-dimensional maps along (p_0, p_1) and (p_0, p_2) .

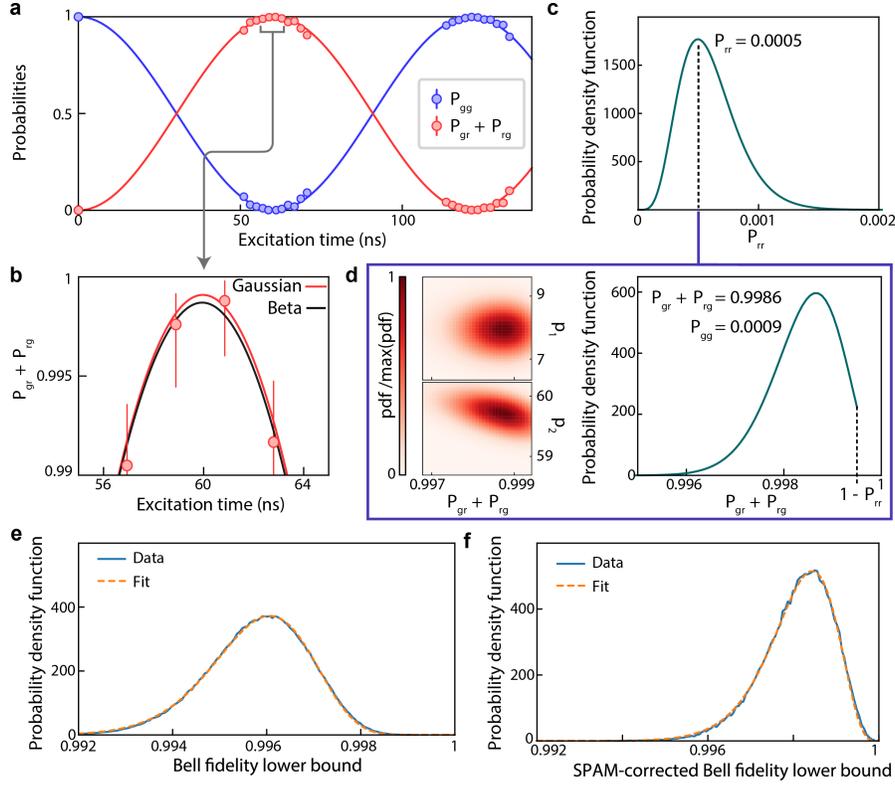


Figure 3.6: **Bell state fidelity measurement from blockaded Rabi oscillations.** **a** Probabilities of measuring both atoms in the ground state P_{gg} and a single atom in the ground state $P_{gr} + P_{rg}$ in the blockaded regime as a function of excitation time. **b**, Zoom over the π time, where we prepare the Bell state $|\Psi^+\rangle$. We use a quadratic fit function of the form $f(x) = p_0 + p_1(x - p_2)^2$ to extract the population values at π and 2π times. **c,d** Fitting method used to obtain the probability distributions of P_{rr} , P_{gg} , and $P_{gr} + P_{rg}$ at π and 2π times. We first experimentally obtain the Beta distribution of the probability P_{rr} to observe both atoms in the Rydberg state. We then perform a joint fit on $P_{gr} + P_{rg}$ and P_{gg} with the same p_1 and p_2 fit coefficients for both. We fix the value of P_{rr} , and condition the joint fit such that the sum of all probabilities is always equals to one. We repeat this process for various values of P_{rr} . The results shown here are for $P_{rr} = 0.0005$, which is the mode of the obtained Beta distribution for P_{rr} . The fitting method uses the true, experimentally measured Beta distribution of each data point. We obtain corresponding probability density functions for each P_{rr} . We perform this method independently for both π and 2π times. **e**, Resulting Bell state fidelity lower bound using the probability density functions of P_{rr} and $P_{gr} + P_{rg}$. We start by randomly drawing from the P_{rr} distribution, then assign the corresponding probability density function of $P_{gr} + P_{rg}$, and draw a value from it. The asymmetry between P_{gr} and P_{rg} is obtained by averaging over each experimental data point, and is assumed to be Gaussian. We repeat this process 1 million times for both π and 2π times. We obtain the corresponding probability density function. **f**, SPAM-corrected Bell state fidelity lower bound distribution, obtained by correcting the probabilities after randomly drawing them from their respective probability density functions.

We then obtain the one dimensional probability density function for p_0 by integrating over p_1 and p_2 (see Fig. 3.6d). This provides the fitted probability density function of P_{gr+rg} , and hence $P_{gg} = 1 - P_{rr} - P_{gr} - P_{rg}$ at π time. We repeat this process for various values of P_{rr} , for both π and 2π times.

At the end of this process, we obtain different probability density functions for each P_{rr} value. The asymmetry between P_{gr} and P_{rg} is obtained by taking the mean of $P_{gr} - P_{rg}$ at π and 2π times. We assume the underlying distribution to be Gaussian, as $P_{gr} - P_{rg}$ is centered on 0, and can be positive or negative with equal probability.

Bell state fidelity

Now that we have the probability density function for all four probabilities at π and 2π times, we move on to the Bell state fidelity extraction. For both π and 2π , we perform a Monte-Carlo sampling of the Beta distribution of P_{rr} which then leads to a joint probability density function for P_{gr+rg} and P_{gg} . We then sample from this, and use Eq. 3.1 to obtain a value for the Bell state fidelity lower bound. We repeat this process 1 million times, and fit the obtained results using a Beta distribution (see Fig. 3.6e). We observe an excellent agreement between the fit and the data, from which we obtain $F_{\text{Bell}} \geq 0.9962(+10/-13)$, where the quoted value is the mode of the distribution, and the error bars represent 68% confidence interval.

We use the same method to obtain the measurement-corrected Bell fidelity, and the state preparation and measurement (SPAM) corrected one. After drawing the probabilities from the probability density functions, we infer the SPAM-corrected probabilities from our known errors, described in detail previously [34]. During this process, there is a finite chance that the sum of probabilities does not sum up to one. This comes from the fact that the probability density functions and the SPAM correction are uncorrelated, an issue which is avoided for raw Bell fidelity extraction thanks to the correlated fit procedure described above. We employ a form of rejection sampling to alleviate this issue by restarting the whole process in the case of such event. We perform this 1 million times, and fit the obtained results using a Beta distribution (see Fig. 3.6f). We observe an excellent agreement between the fit and the data, from which we obtain a SPAM-corrected fidelity $F_{\text{Bell}} \geq 0.9985(+7/-12)$, where the quoted value is the mode of the distribution, and the error bars represent 68% confidence interval.

3.3.5 Interaction limitation for Bell fidelity

For the case of measuring two-qubit Bell state fidelities we set $\Omega/2\pi = 6.2$ MHz. Interaction strengths in Fig. 3.4b are directly measured at inter-atomic separations of 4 and 5 μm , and extrapolated via the predicted $1/r^6$ scaling to the level at 2.5 μm . Mean atomic distances are calibrated via a laser-derived ruler based on shifting atoms in coherent superposition states [28]. We calibrate $C_6/2\pi = 230(25)$ GHz μm^6 using maximum-likelihood-estimation (and associated uncertainty) from resonant quench dynamics [29], which additionally calibrates a systematic offset in our global detuning.

We then estimate the theoretically expected Bell state fidelity using perturbation analysis. Specifically, resonant blockaded Rabi oscillation are described by

$$\hat{H}/\hbar = \frac{\Omega}{2}(\hat{X}_1 + \hat{X}_2) + V\hat{n}_1\hat{n}_2, \quad (3.6)$$

where $V = C_6/r^6$ is the distance-dependent, interaction strength between two atoms separated at distance r (see Eq. 2.1). Since the two-atom initial ground state, $|\psi(0)\rangle = |gg\rangle$, has even parity under the left-right reflection symmetry, the Rabi oscillation dynamics can be effectively solved in an even-parity subspace with three basis states of $|gg\rangle$, $|rr\rangle$, and $|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|gr\rangle + |rg\rangle)$. In the Rydberg-blockaded regime where $V \gg \Omega$, we can perform perturbation analysis with the perturbation parameter $\eta = \Omega/\sqrt{2}V$ and find that the energy eigenvectors of the subspace are approximated as

$$\begin{aligned} |E_1\rangle &\approx \frac{(1 - \frac{\eta}{4} - \frac{\eta^2}{32})|gg\rangle + (-1 - \frac{\eta}{4} + \frac{17\eta^2}{32})|\Psi^+\rangle + (\eta - \frac{3\eta^2}{4})|rr\rangle}{\sqrt{2}} \\ |E_2\rangle &\approx \frac{(-1 - \frac{\eta}{4} + \frac{\eta^2}{32})|gg\rangle + (-1 + \frac{\eta}{4} + \frac{17\eta^2}{32})|\Psi^+\rangle + (\eta + \frac{3\eta^2}{4})|rr\rangle}{\sqrt{2}} \\ |E_3\rangle &\approx \eta^2|gg\rangle + \eta|\Psi^+\rangle + (1 - \eta^2/2)|rr\rangle \end{aligned}$$

with their corresponding energy eigenvalues of $E_1 \approx V(-\eta - \eta^2/2)$, $E_2 \approx V(\eta - \eta^2/2)$, and $E_3 \approx V(1 + \eta^2)$, respectively¹. Rewriting the initial state using the perturbed eigenbasis, we solve

$$F_{\text{Bell}} = \max_t |\langle \Psi^+ | e^{-i\hat{H}t} | \psi(0) \rangle|^2 \quad (3.7)$$

to obtain the analytical expression of the maximum achievable Bell state fidelity, F_{Bell} , at a given perturbation strength η . Keeping the solution up to the second order of η , we find

$$F_{\text{Bell}} = 1 - \frac{5}{4}\eta^2 = 1 - \frac{5}{8} \left(\frac{\Omega}{V} \right)^2 \quad (3.8)$$

obtained at $t = \pi/\sqrt{2}\Omega$.

¹Note that the definition of E_3 (both the eigenvalue and eigenvector) is slightly corrected here from the version presented in Ref. [26].

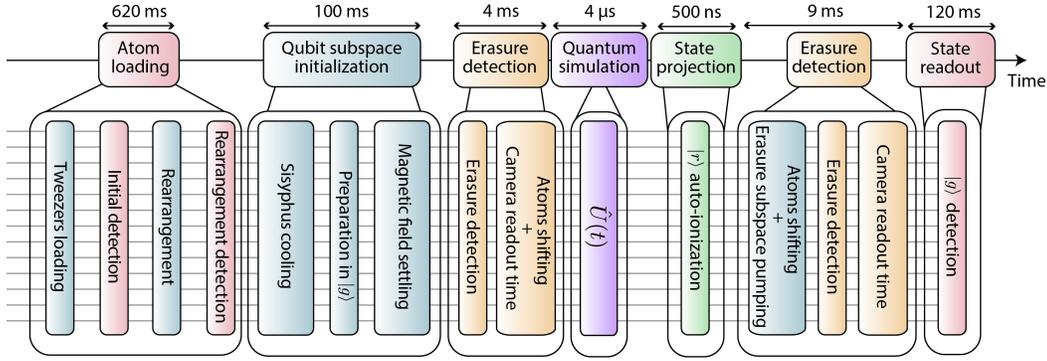


Figure 3.7: **Experimental sequence.** Sketch of the experimental sequence including the erasure detection for the GHZ adiabatic preparation experiment. We start by loading the atoms into the desired geometry, then initialize the atoms in $|g\rangle$, perform the Bell state generation or quantum simulation, and finally read out by auto-ionizing atoms in $|r\rangle$ and imaging atoms in $|g\rangle$. We perform two erasure images: one before applying $\hat{U}(t)$, and one after auto-ionization.

3.4 CASE STUDY: GHZ STATE PREPARATION

Having demonstrated the benefits of erasure-excision for the case of improving two-qubit entanglement fidelities, we now show it can be similarly applied to the case of many-body quantum simulation, demonstrating the utility of erasure detection for NISQ applications. As part of this investigation, we distinguish erasure errors from preparation and Rydberg spontaneous decay, the latter of which becomes more visible in a many-body setting and for longer evolution times. Further, we show that erasure conversion can reveal otherwise hidden correlations between errors.

3.4.1 Overview

As a prototypical example, we explore a quasi-adiabatic sweep into a \mathbb{Z}_2 -ordered phase (Fig. 3.8a) through the use of a varying global detuning [156] (Fig. 3.8b). In this ordered phase, ground and Rydberg states form an antiferromagnetic (AFM) pattern, with long-range order appearing at a quantum phase transition. Unlike previous examples [41, 157], we operate in the effectively *attractive* interacting regime of the Rydberg blocked space [156], which features a true two-fold degenerate ground state for systems with an even number of atoms, even for open boundary conditions, and without explicitly modifying the boundary [157].

For performing many-body quasi-adiabatic sweeps, the detuning is swept symmetrically in

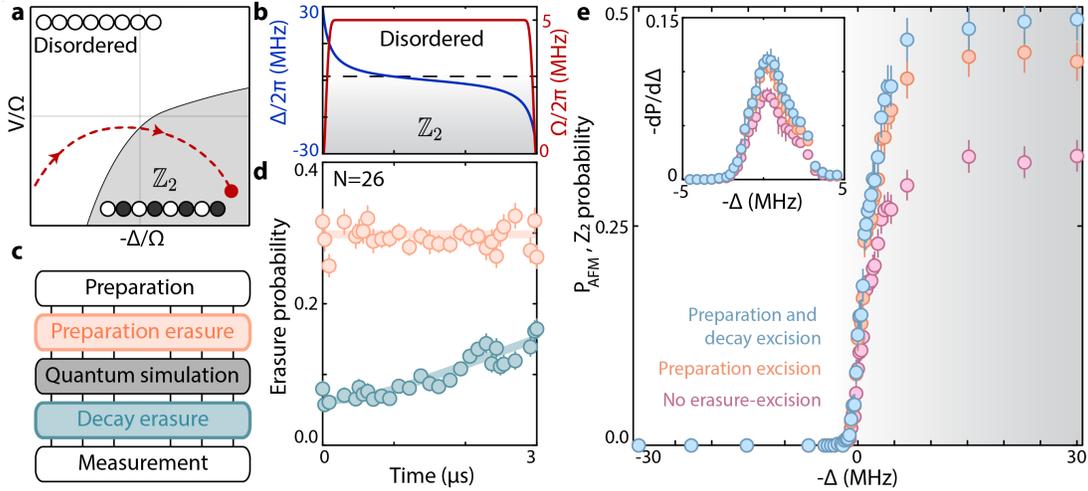


Figure 3.8: **Erasure conversion in quantum simulation.** **a, b,** We perform quasi-adiabatic sweeps with $N = 26$ spins in the effective ground state manifold of an attractive Rydberg Hamiltonian, starting from the initially disordered phase, and ending in the \mathbb{Z}_2 ordered phase. **c,** We perform two erasure images, one preceding the quantum simulation (to detect preparation errors), and one following (to partially detect Rydberg decay errors). **d,** The probability for detecting a preparation error (orange markers) remains constant while the probability for detecting a decay error (green markers) grows in agreement with the Rydberg lifetime and detection infidelities (solid lines). **e,** The total probability, P_{AFM} , for forming either of the AFM \mathbb{Z}_2 states is improved by performing erasure-excision on all errors (blue markers), as compared to only on preparation errors (orange markers) or performing no excision (pink markers). The sensitivity of P_{AFM} with respect to a change in Δ also increases with erasure-excision (inset).

a tangent profile from +30 MHz to -30 MHz, while the Rabi frequency is smoothly turned on and off with a maximum value of $\Omega/2\pi = 5.6$ MHz. For an initially positive detuning, the $|r\rangle$ state is energetically favorable, making the all ground initial state, $|gg\dots gg\rangle$, the highest energy eigenstate of the blockaded energy sector, where no neighboring Rydberg excitations are allowed. For negative detunings, where $|g\rangle$ is energetically favorable, the highest energy state uniquely becomes the symmetric AFM state $(|grgr\dots gr\rangle + |rgrg\dots rg\rangle)/\sqrt{2}$ in the deeply ordered limit. Thus, considering only the blockaded energy sector, sweeping the detuning from positive to negative (remaining in the highest energy eigenstate) is equivalent to the ground state physics of an effective Hamiltonian with attractive Rydberg interaction and inverted sign of the detuning. This equivalence allows us to operate in the effectively attractive regime of the blockaded phase diagram of Ref. [156]. For our Hamiltonian parameters, we use exact diagonalization numerics to identify the infinite-size critical detuning using a scaling collapse near the finite-system size minimum energy gap [158].

Staying adiabatic during ground state preparation requires evolution over microseconds, orders of magnitude longer than the two-qubit entanglement operation shown before, which

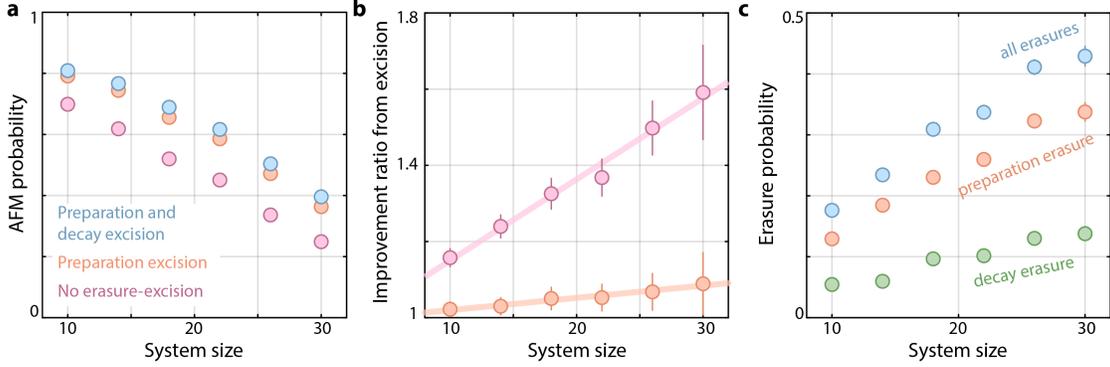


Figure 3.9: System size scaling of sweep fidelity. **a**, Total probability for forming either of the \mathbb{Z}_2 states following the quasi-adiabatic sweeps presented in Fig. 3.8 of the main text, here presented for the final sweep time as a function of system size for full erasure-excision (blue markers), preparation erasure-excision (orange markers), and the baseline (pink markers) data. **b**, The ratio gain from using full erasure-excision grows as a function of system size, both with respect to the baseline values (pink markers), and to the case of only excising preparation erasures (orange markers). Solid lines are linear guides to the eyes. **c**, For a fixed detection threshold of 5 photons, the number of erasure errors also increases as a function of system size.

magnifies the effect of Rydberg decay. In order to differentiate between leakage out of the qubit manifold due to either preparation errors or Rydberg decay, we perform two erasure images, one before the adiabatic sweep which captures preparation errors, and one after (Fig. 3.8c). The second image allows us to measure Rydberg decay into the detection subspace throughout the sweep. For a system size of $N = 26$ atoms (Fig. 3.8d), we see the number of detected preparation erasures (orange markers) stays constant over the course of a $3 \mu\text{s}$ sweep; conversely, the number of detected decay erasures (green markers) grows over time, in good agreement with the measured Rydberg lifetime and erasure image infidelities (green solid line).

With the ability to distinguish these effects, we plot the total probability to form either of the AFM states, $P_{\text{AFM}} = P(|grgr\dots gr\rangle) + P(|rgrg\dots rg\rangle)$ (Fig. 3.8e). At the conclusion of the sweep, we find $P_{\text{AFM}} = 0.33(2)$ without any erasure-excision (pink markers). By excising instances with preparation erasures, this fidelity is improved to $0.44(2)$ (orange markers), and is then further improved to $0.49(2)$ by additionally excising Rydberg decay erasures. The sharpness of the signal, exemplified by the derivative of P_{AFM} with respect to the detuning, is similarly improved near the phase boundary (Fig. 3.8e, inset). We also observe that the gain in P_{AFM} from erasure-excision increases with system size (Fig. 3.9).

We further explore how errors affect quantities reflecting higher-order magnetization statis-

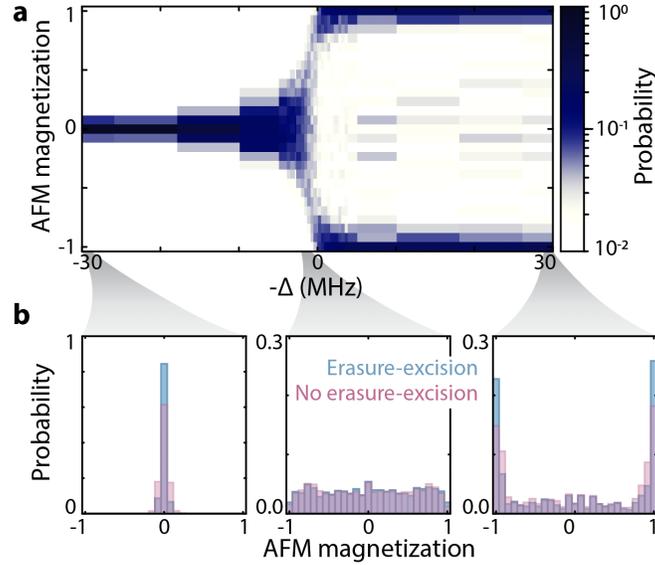


Figure 3.10: **Erasure conversion in quantum simulation.** **a**, The probability distribution for measuring a given AFM magnetization after the adiabatic GHZ-preparation sweep is initially peaked at 0 in the disordered phase, before bifurcating when entering the \mathbb{Z}_2 phase, consistent with spontaneous symmetry breaking. **b**, Deep in either phase, erasure-excision leads to a sharpening of the probability distribution (left, right). Around the phase transition, we observe a close-to-flat distribution (middle).

tics [159]. To this end, we explore the probability distribution to find magnetic order of different magnitude by studying the AFM magnetization operator, defined as

$$\hat{M} = \hat{Z}_A/N_A - \hat{Z}_B/N_B, \quad (3.9)$$

where $\hat{Z}_S = \sum_{j \in S} \hat{Z}_j$ is the total magnetization operator in sub-lattice $S=A$ (odd sites) or $S=B$ (even sites), respectively, N_S is the number of atoms in each sub-lattice, and $\hat{Z}_j = |r\rangle\langle r| - |g\rangle\langle g|$ is the local magnetization at site j . We plot the probability to find a specific eigenvalue, M , of \hat{M} as a function of detuning (Fig. 3.10a). While the values of M are initially tightly grouped around $M = 0$ in the disordered phase, as the sweep progresses the probability distribution bifurcates, forming two separate dominant peaks in the \mathbb{Z}_2 phase, consistent with aforementioned two-fold spontaneous symmetry breaking across the quantum phase transition. We find that erasure-excision improves the sharpness of the distribution in both the disordered and \mathbb{Z}_2 phases (Fig. 3.10b). Near the phase transition, the distribution is close-to-flat, consistent with order appearing at all length scales. We note, though do not explore in more detail, that this observation is again consistent with the emergence of microscopic critical phenomena [159–161].

These results demonstrate improvements in fidelity for preparation of long-range-ordered ground states with erasure-excision in quantum simulation experiments, a first proof-of-

principle for utilizing erasure conversion in NISQ-type applications. Our demonstrated improvements for analog quantum simulation of ground-state physics could be extended to non-equilibrium dynamics, for example generating large entanglement entropy states [29], with the potential to reach a quantum advantage over classical simulations [162]. We envision erasure conversion will improve other tasks such as quantum optimization [163] and potentially quantum metrology [164, 165].

3.4.2 Learning from erasure errors

Finally, we turn to studying a new tool enabled by our implementation of erasure conversion: exploring the effect of errors on experimental outcomes at a microscopic level and studying correlations between different error sources, which is enabled by having three separate images for a given experimental run (Fig. 3.11a). In particular, we consider the joint probability distribution, $\mathcal{P}(e_1^{(i)}, e_2^{(j)}, e_3^{(k)})$, that atoms at sites i , j , and k are detected, respectively, in the preparation erasure image (e_1), the decay erasure image (e_2), and the final state detection image (e_3).

We again consider adiabatic sweeps into the \mathbb{Z}_2 phase as in Fig. 3.8, but now with a total duration of $8 \mu\text{s}$. We first study $\mathcal{P}(e_3^{(j+d)} = 0 | e_1^{(j)} = 0)$, equivalent to finding a Rydberg excitation on site $j + d$, conditioned on finding no preparation erasure on site j . We plot this quantity (Fig. 3.11b, left) as a function of both d and the sweep duration. We explicitly average over choices of j and find a signal essentially uniform in d .

However, if we instead consider $\mathcal{P}(e_3^{(j+d)} = 0 | e_1^{(j)} = 1)$, the probability to find a Rydberg excitation on site $j + d$ conditioned on *detecting* a preparation erasure on site j , markedly different behavior emerges (Fig. 3.11b, middle). For simplicity, we further post-select on instances where only a single erasure is detected across the entire array. At intermediate sweep times, we observe an AFM order forms around the preparation erasure error position. We interpret the error as breaking the atom chain into two shorter chains; excitations will naturally form at the system edges of these shorter chains in order to maximize the Rydberg density in the attractive regime in which we operate. This effectively pins the Rydberg density around the error, which then establishes a preferred AFM order further out into the array. Interestingly, the equivalent quantity for decay erasures, $\mathcal{P}(e_3^{(j+d)} = 0 | e_2^{(j)} = 1)$, shows a more complex behavior.

To quantify this behavior more explicitly, we consider a variant of the AFM magnetization (Eq. 3.9) conditioned on the erasure location, where sublattice A (B) is now defined as being sites an odd (even) distance away from an erasure. In Fig. 3.11c we plot the mean AFM

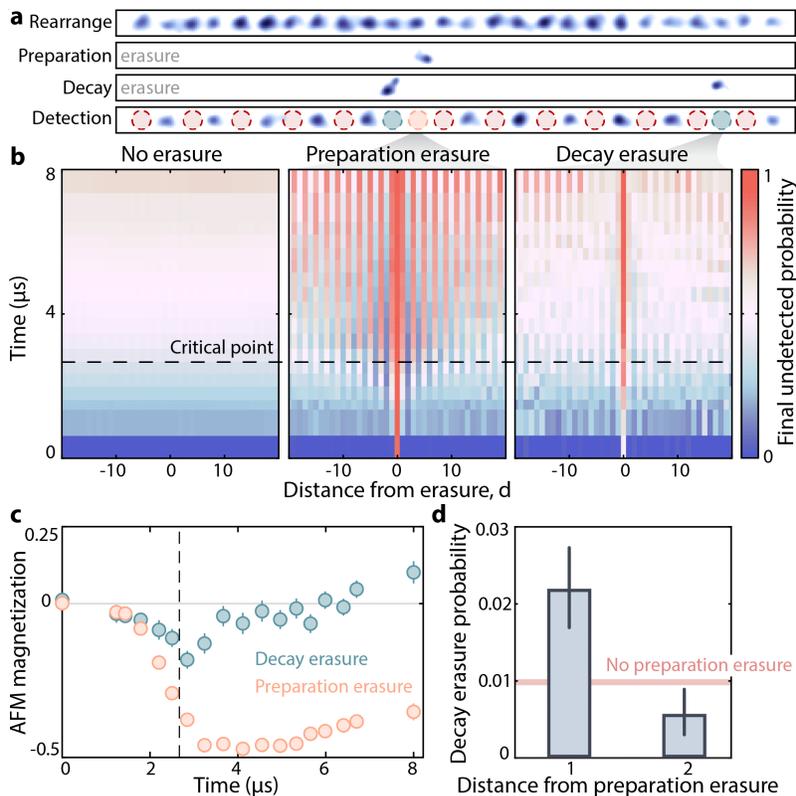


Figure 3.11: **Learning from erasure errors.** **a**, Post-processed [138] single-shot atom fluorescence images. After arranging the array (top panel) we use the fast erasure images to learn how detected errors (middle panels) affect detection outcomes (lower panel). **b**, Conditional probability to detect no atom in the final image as a function of sweep progress and distance from a hypothetical erasure event. In the case of no erasures (left), the resulting profile is uniform. However, when conditioning on detecting a preparation erasure (middle), the error breaks the \mathbb{Z}_2 symmetry by establishing a single AFM order. In the case of conditioning on decay errors (right) the situation is more complex. **c**, AFM magnetization (see Eq. 3.9) as a function of time. Preparation erasures (orange markers) lead to a growth of a single AFM order with Rydberg excitations predominantly on sub-lattice A (defined as sites an odd distance from the erasure position). Decay erasures (green markers) follow a similar trend at early times by acting as effective preparation errors, but past the critical point (vertical dashed line) their behavior reverses: a decay erasure spontaneously breaks the two-fold symmetry, where the neighboring sites must have been in the ground state due to the \mathbb{Z}_2 structure, yielding Rydberg excitations on sub-lattice B . **d**, For the maximum sweep duration, the probability of detecting a Rydberg decay erasure (bars) is significantly increased (decreased) from the baseline level (red fill) at a distance of one (two) away from the site of a detected preparation erasure, induced by the altered Rydberg population on these sites due to the crystal formation in **b**, indicating direct detection of correlations between errors through erasure imaging.

magnetization for both the preparation (orange circles) and decay erasure (green circles) cases. Preparation erasures develop a negative, single AFM order as they pin Rydberg

excitations at odd distances away from the erasure.

Decay erasures behave similarly before the critical point, as Rydberg decay acts effectively as a preparation error. However, past the critical point, this behavior changes: decay now acts as a measurement on the AFM superposition ground state, selecting one of these orders. In this case, assuming perfect \mathbb{Z}_2 states, the neighboring sites must have been in the ground state to detect a decay, meaning the AFM order is reversed from the preparation case. This leads to data that first dips to negative values and then grows to positive values past the phase transition (green markers in Fig. 3.11c).

We also study correlations between preparation errors and Rydberg decay. In particular, a preparation error forces atoms at odd intervals from the preparation erasure to have a higher probability to be in Rydberg states, meaning they should also be more likely to decay. As shown in Fig. 3.11d, we directly observe this effect at the end of the sweep by considering $\mathcal{P}(e_2^{(j+d)} = 1 | e_1^{(j)} = 1)$, the probability to detect a decay erasure at a distance d away from a preparation erasure. For $d = 1$ ($d = 2$), this probability is significantly increased (decreased) from the unconditional decay erasure probability, in line with the increased (decreased) Rydberg population on these sites, which shows that errors are correlated.

We note that erasure-excision for preparation errors using the first erasure image can be considered heralding the subsequent quantum simulation on the presence of atoms in tweezers in the correct initial state. For erasure-excision of Rydberg decay using the second erasure image, we interpret the post-selected results as coming from a non-jump trajectory in a Monte Carlo wavefunction approach [166].

Overall, insights into erasure-error correlations could be used to understand error processes in NISQ devices in unprecedented detail, in particular if erasure detection could be made time-resolved with respect to the many-body dynamics. This could also be used to study post-measurement physics with erasure detection, such as measurement-induced phase transitions [167, 168] and measurement-altered quantum criticality [169].

3.4.3 Statistics reduction due to erasure-excision

Our demonstration of erasure-excision explicitly discards some experimental realizations (Fig. 3.9), which can be seen as a downside of the method. However, this is a controllable trade-off: by adjusting the threshold for detecting an erasure error, we can balance gains in fidelity versus losses in experimental statistics (as shown in Fig. 3.12) for whatever particular task is of interest. In general, the optimum likely always includes some amount of erasure-excision, as it is usually better to remove erroneous data than keeping them.

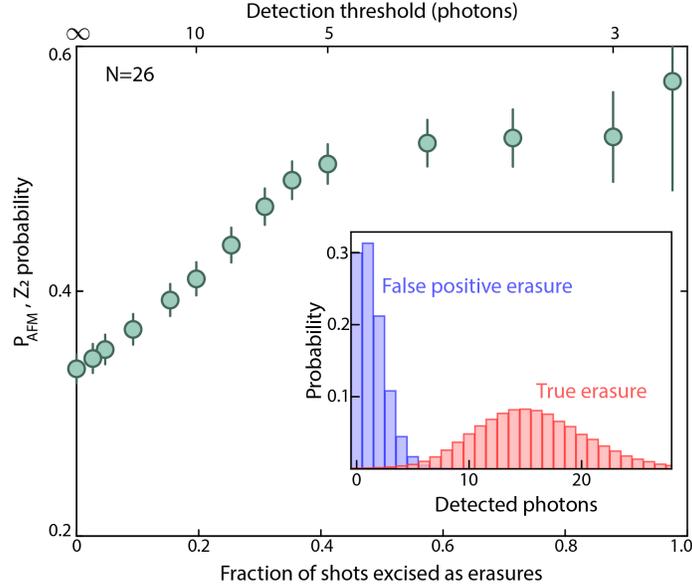


Figure 3.12: **Controlling fidelity gains from erasure-excision.** Erasure-excision explicitly trades improved experimental fidelity for a reduced number of experimental repetitions. However, this process is controllable by adjusting the threshold used for detecting atoms during erasure images. Changing the threshold essentially changes the false positive and false negative rate for detecting erasures correctly (inset). We plot the total AFM probability (green markers) after a sweep as in Fig. 3.8 of the main text, and vary the detection threshold used for identifying erasures. For too high a threshold, many erasure events go unnoticed, and so erroneous outcomes become relatively more prevalent, reducing the overall fidelity. As the detection threshold is lowered, more true erasures, where an atom is actually present in the erasure image, are correctly detected, which improves the fidelity. However, lowering the threshold too far (in our case past ~ 5 photons) increases the likelihood of seeing false positive erasures; excising data based on these events discards experimental statistics with relatively little gain in fidelity. In the main text, we select a detection threshold of 5 photons.

For erasure detection employed in the context of many-body quantum simulation, we adjust the binarization threshold for atom detection to raise the false positive imaging fidelity to 0.9975, while the false negative imaging fidelity is lowered to ~ 0.6 (Fig. 3.8d); this is done as a conservative measure to prioritize maximizing the number of usable shots while potentially forgoing some fidelity gains (Fig. 3.12).

3.5 ERASURE CORRECTION COOLING

Next, we demonstrate that erasure conversion is not solely limited to improving quantum simulation fidelities through detecting leakage errors. Specifically, we showcase a new measurement-based [170] cooling method which uses erasure conversion (and correction)

to prepare atoms nearly deterministically in their motional ground state – we term our method *erasure correction cooling*.

This technique is species-agnostic and can be used to efficiently prepare motional ground state atoms. While cold atoms are important for practically all applications, later in this thesis (Chapter 4) we will discuss how high-fidelity motional purification as we perform here enables new tasks in quantum information processing.

The mechanism behind erasure cooling is fundamentally reminiscent of the Maxwell’s demon thought experiment, and relies on (i) conversion of motional excitations into detectable, localized erasure errors [170], and (ii) their active correction. Using such a technique, we obtain a motional ground state occupation of $0.984(+4/-5)$, increased to $0.995(+3/-4)$ if post-selection is used rather than active correction. Importantly, we find this new cooling scheme systematically outperforms traditional sideband cooling, which we demonstrate experimentally in the case of shallow tweezers by comparing our results to simulations of an idealized sideband cooling process.

3.5.1 Overview

To start, we describe erasure correction cooling (ECC). The basic principle relies on detecting atoms in motion-excited states, while shelving motion-ground state atoms such that they are not disturbed by measurement. From there, motional excitations can either be removed, or selectively re-cooled. The results presented here concern the tightly-confined, radial direction of the optical tweezers; we also implemented ECC along the weakly-confining, axial direction, for which we find similar results in Sec. 3.5.6.

Atoms are initialized in their ground electronic state 1S_0 and initially cooled with Sisyphus cooling [85] (a conventional form of laser cooling) to an initial motional ground state population of $P'_0 = 0.77(1)$. To perform ECC, we convert remaining motional excitations into erasure errors (Fig. 3.13a) by working in the *sideband resolved* regime on the transition between 1S_0 and the metastable excited state 3P_0 . The trapping wavelength is 813 nm to achieve a magic trapping condition for 1S_0 and 3P_0 .

We note that in all of the following, we will consider a single motional degree of freedom, which is defined along the $^1S_0 \leftrightarrow ^3P_0$ laser beam propagation axis. We use numerical notation for motional states (e.g. $|0\rangle$ for the motional ground state), and pseudo-spin notation for electronic states (e.g. $|\downarrow\rangle$ is 1S_0 and $|\uparrow\rangle$ is 3P_0); for the latter we will use the terms spin, electronic state, and optical qubit interchangeably. We further denote their tensor product as e.g. $|0\rangle \otimes |\uparrow\rangle = |0, \uparrow\rangle$.

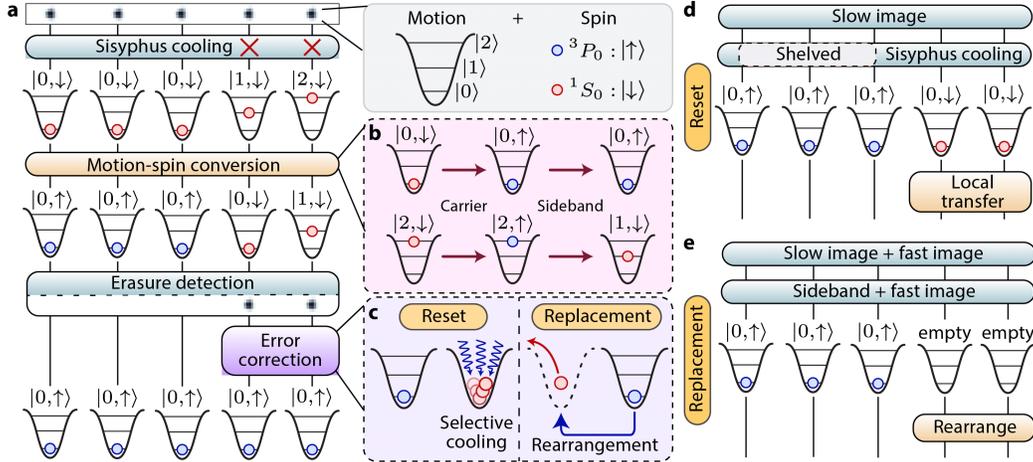


Figure 3.13: **Improving motional ground state preparation with erasure correction cooling.** **a,b**, After cooling the atoms with Sisyphus cooling, we apply spin-motion conversion, which leaves atoms in a different pseudo-spin state depending on whether they do or do not have motional excitations. To accomplish spin-motion conversion, we transfer the atoms in $|\downarrow\rangle$ to the $|\uparrow\rangle$ state using a carrier pulse, which conserves the motional state of the atoms. We then apply a blue sideband pulse, which performs the $|n, \uparrow\rangle \rightarrow |n-1, \downarrow\rangle$ transition, where n is the atom's motional level. We then detect atoms with motional excitations by imaging the pseudo-spin state $|\downarrow\rangle$, without perturbing atoms in the motional ground state, which reside in $|\uparrow\rangle$. **c**, We correct tweezers containing atoms with motional excitations in two different ways: (i) by reapplying Sisyphus cooling selectively on such atoms (reset), or (ii) by replacing such atoms using rearrangement (replacement). **d**, In the reset case, the erasure detection image is a slow and non-destructive image. After the image, we cool the atoms again. Atoms in $|0, \uparrow\rangle$ are shelved from the Sisyphus cooling mechanism. We then selectively transfer [33] the atoms which were imaged in $|\downarrow\rangle$ to $|\uparrow\rangle$. **e**, For the replacement method, the erasure detection image is a combination of a slow, high survival image, and a fast, low survival image. We then repeat a sideband pulse and a fast image in order to improve the erasure cooling efficiency. We then fill the emptied tweezers with auxiliary atoms which were measured to be in the $|0, \uparrow\rangle$ state.

Erasure correction cooling then proceeds with the following steps (Fig. 3.13a,b): first we drive all atoms into the excited electronic state using the motion-preserving *carrier* transition, $|n, \downarrow\rangle \rightarrow |n, \uparrow\rangle$, where n is an arbitrary motional level. Second, we perform a spin-motion coupling *sideband* transition, $|n, \uparrow\rangle \leftrightarrow |n-1, \downarrow\rangle$, ideally leaving only motional ground state atoms in the excited electronic state. We then perform fluorescence imaging of the ground electronic state, while leaving atoms in $|0, \uparrow\rangle$ unperturbed [26]. This means that we directly image the location of motion-excited atoms, for which we can then perform error correction.

This correction can be done in two ways (Fig. 3.13c-e), either: (1) *Reset*, where we reapply Sisyphus cooling with site-selectivity, or (2) *Replacement*, by discarding hot atoms and

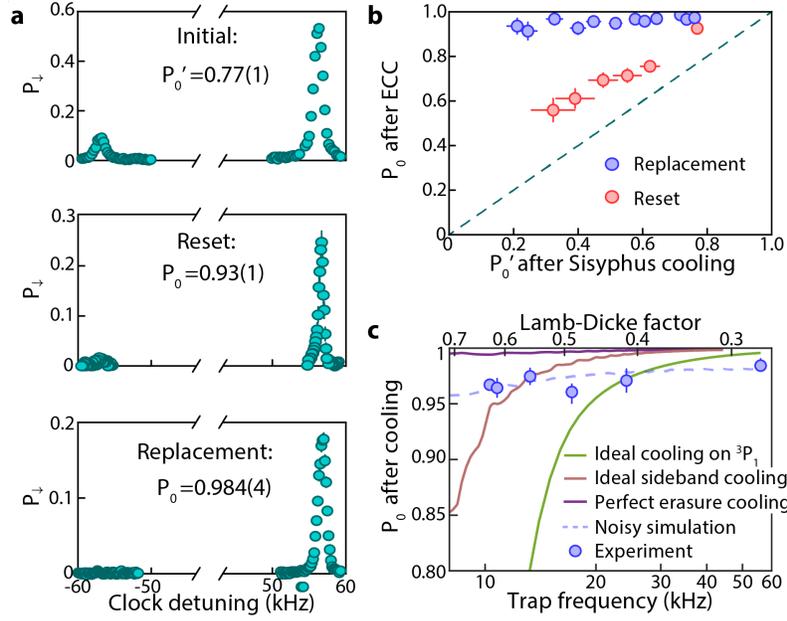


Figure 3.14: **Erasure correction cooling in the radial direction** **a**, Probability, P_0 , for an atom to be in the motional ground state, $|0\rangle$, measured via sideband spectroscopy. We show the results without erasure correction cooling, performing erasure correction with reset, and erasure correction with replacement. **b**, Post-correction motion-ground state probability, P_0 , as a function of intentionally varied motion-ground state probability after Sisyphus cooling, P'_0 . We show the results for both replacement (blue markers) and reset (red markers). The replacement method is largely insensitive to the Sisyphus cooling efficiency. The green dashed line is a guide to the eye which represents unity gain. **c**, P_0 as a function of the trap frequency ω . Our experimental erasure cooling results (blue markers) using the replacement method are in good agreement with an *ab initio* error model including known imperfections (blue dashed line). We compare against two variants of sideband cooling, one using the 3P_1 state of strontium (green line), and one which is an idealization representing the ultimate limit of sideband cooling for given trap and Rabi frequencies, and Lamb-Dicke factor. Our experimental erasure cooling outperforms both sideband cooling variants in the regime of low trap frequency, and an error-free simulation of erasure cooling is superior to sideband cooling for all frequencies. Error bars represent the standard error of the mean.

replacing them with cold atoms using atom rearrangement. At the end of the correction, we transfer any remaining atom in $|\downarrow\rangle$ to $|\uparrow\rangle$ using optical pumping to avoid any vacancy defects in the array.

To assess the efficacy of ECC, we use a 39-tweezer array, and measure P_0 via sideband spectroscopy on the $|\downarrow\rangle \rightarrow |\uparrow\rangle$ transition (Fig. 3.14a). Starting from $P'_0 = 0.77(1)$ after Sisyphus cooling, we obtain a final motional ground state population $P_0 = 0.93(1)$ with one round of reset, and $P_0 = 0.984(+4/-5)$ with replacement. The maximum P_0 after reset is limited here by only doing one round of cooling, with a theoretical maximum of $P_0 = P'_0 + (1 - P'_0)P'_0 \approx 0.95$. In Fig. 3.14b we show the efficacy of ECC as a function of

P'_0 obtained after Sisyphus cooling (allowing up to 3 rounds of reset cooling for the lowest P'_0 values); in general we find that ECC with replacement is especially insensitive to this initial cooling temperature. If we post-select on erasure detection, rather than applying active correction, we achieve $P_0 = 0.995(+3/-4)$ (not shown).

Having demonstrated the efficacy of ECC for our particular experimental realization, we now show that it outperforms even idealized, noise-free, conventional laser cooling in certain realistic parameter regimes (Fig. 3.14c). We first perform a noise-free simulation of ECC (purple line), and compare it against two possible realizations of ideal sideband cooling. The first is against cooling on the commonly used $^1S_0 \rightarrow ^3P_1$ intercombination line [99, 100] (green line), while the second is against idealized sideband cooling performed with infinitely small linewidth (red line), but including an effective spontaneous decay in order to remove entropy from the system [171].

We observe that across all trap frequencies, the idealized ECC always outperforms idealized sideband cooling, but particularly so in the limit of low trap depth, where the sideband-resolved condition ($\omega \gg \Gamma$) breaks down. In this limit, sideband cooling efficiency becomes limited by (i) the finite Rabi frequency (here $2\pi \times 2.5$ kHz) used to cool the atoms, and (ii) photon recoil. Conversely, ECC is (i) less sensitive to errors coming from the finite Rabi frequency because they can be partially erased, and (ii) not limited by photon recoil since cooling does not rely on spontaneous emission from a cooling state. Instead, ECC is primarily limited by tweezer-induced heating during error correction, as seen in both experimental (blue points) and *ab initio* error model results (dashed line) accounting for known noise sources. Importantly, for $\omega/(2\pi) \lesssim 24$ kHz, the experimental ECC implementation outperforms sideband cooling on $^1S_0 \rightarrow ^3P_1$, and outperforms even the idealized sideband cooling for $\omega/(2\pi) \lesssim 12$ kHz.

We stress that this regime of relatively low trap frequency is particularly relevant for the tweezers' weakly confining, axial direction. Furthermore, whilst a high trap frequency enables large sideband-cooling efficiency, it also limits the achievable number of physical qubits for a fixed laser power. Our results demonstrate that the same cooling can be achieved practically at a much lower trap depth, potentially opening the door to greatly expanding the size of tweezer arrays.

Our scheme for erasure correction cooling (i) presents a practical advantage in cooling efficiency for our setup, and (ii) can outperform sideband cooling in general, which we demonstrated experimentally in the regime of low trap frequency.

3.5.2 The mechanism of erasure correction cooling

Here we detail the erasure correction cooling (ECC) mechanism in greater detail. We start with the erasure detection of motional excitations, then describe the two methods employed to perform error correction.

The erasure detection is described in Fig. 3.13. After cooling the atoms with Sisyphus cooling using the 3P_1 level, we transfer the atoms in $|\downarrow\rangle$ to the $|\uparrow\rangle$ state using a carrier pulse, which conserves the motional state of the atoms. We then apply a blue sideband pulse which performs the $|n, \uparrow\rangle \rightarrow |n - 1, \downarrow\rangle$ transition, where n is the atom's motional level. Atoms in $|0, \uparrow\rangle$ are not transferred back to $|\downarrow\rangle$. We then selectively image atoms in $|\downarrow\rangle$, which allows us to spatially resolve the position of motional excitations in the array. We then correct for these errors using two methods: re-using the same atoms by selectively cooling them again (reset) or replacing them with auxiliary atoms in the $|0, \uparrow\rangle$ state (replacement).

We first detail the reset method (Fig. 3.13d). In this case, the erasure detection is a slow, high survival image with a duration of 45 ms [85]. After the image, we cool the atoms in $|\downarrow\rangle$ again through Sisyphus cooling, while atoms in $|0, \uparrow\rangle$ are shelved from this cooling. We then use movement-induced phase shifts [33] (described in Chapter 4) to selectively transfer to $|\uparrow\rangle$ the atoms which were previously imaged in $|\downarrow\rangle$. This process can be repeated several times in order to improve the erasure cooling efficiency.

The erasure detection image in the case of the replacement method (Fig. 3.13e) is a combination of a high survival, slow image [85] with a duration of 30 ms and a fast, low survival [26] image with a duration of $24 \mu\text{s}$. We perform this combination in order to maximize the imaging fidelity while minimizing the misidentification of Raman scattering events from the $|\uparrow\rangle$ state to the $|\downarrow\rangle$ state. We then perform another blue sideband pulse and a fast image in order to minimize the effect of tweezer-induced heating during the erasure detection, which improves the eventual erasure cooling efficiency. Due to the fast image, the tweezers holding atoms in $|\downarrow\rangle$ are empty. We fill them with auxiliary atoms which were measured to be in the $|0, \uparrow\rangle$ state.

We ramp down the tweezers depth to $\sim 40 \mu\text{K}$, which corresponds to one-tenth of their nominal depth in order to reduce the effect of tweezer-induced heating such as Rayleigh scattering. In particular, we also observe that rearrangement-induced heating strongly depends on the tweezers' depth, and is minimized for a tweezer depth in the range $\sim 20 - 40 \mu\text{K}$. This rearrangement-induced heating is identified to be coming from (i) the finite probability that an atom is heated during transportation, and (ii) tweezers depth oscillations during atom movement due to intermodulation of the varying frequencies in the acousto-

optic deflector used to generate the tweezers [81], although we leave a quantitative analysis of these effects to future work.

3.5.3 Cooling simulations

Here we provide details on the cooling simulations results presented in Fig. 3.14c. We first describe sideband cooling on the 3P_1 level of strontium, then the idealized sideband cooling, and finally both idealized and noisy erasure cooling using replacement.

In all cases, we simulate the following one-dimensional Hamiltonian

$$H = \hbar\omega(a^\dagger a + \frac{1}{2}) - \hbar\Delta|\uparrow\rangle\langle\uparrow| + \frac{1}{2}\hbar\Omega(e^{i\eta(a+a^\dagger)}|\uparrow\rangle\langle\downarrow| + \text{h.c.}), \quad (3.10)$$

where a^\dagger (a) is the creation (annihilation) operator acting on the motional levels, ω is the trapping frequency, Δ is the detuning, η is the Lamb-Dicke factor and Ω is the Rabi frequency. We set $\Delta = -\omega$, and use a basis of ten motional states.

Ideal sideband cooling on 3P_1

We perform a simulation of sideband cooling on the $|^3P_1, m_J = 0\rangle$ state. Here we assume ideal conditions, namely: (i) we assume no laser noise, (ii) we set a Rabi frequency 500 times smaller than the cooling state linewidth, and (iii) we consider cooling only in a single direction, but take into account the dipole radiation pattern in 3D (more details below).

We solve the Master equation in which we take into account spontaneous emission from the $|\uparrow\rangle$ state and its associated momentum kick which follows a dipole pattern distribution:

$$i\hbar\partial_t\rho = (H_\Gamma\rho - \rho H_\Gamma^\dagger) + i\hbar\Gamma \int d\theta D(\theta)c_\theta\rho c_\theta^\dagger.$$

We define $H_\Gamma = H - \frac{i}{2}\hbar\Gamma|\uparrow\rangle\langle\uparrow|$, and $c_\theta = e^{-i|k|\cos(\theta)}|\downarrow\rangle\langle\uparrow|$. Here, ρ is the density matrix of the spin and motional states, $\Gamma = 2\pi \times 7.4\text{kHz}$ is the linewidth of 3P_1 , $D(\theta) = 3\sin^2(\theta)/(8\pi)$ is the dipole radiation pattern for the transition of interest, θ the angle of emission with respect to the quantization axis. The dipole emission pattern takes this simple form thanks to considering cooling on the $|^3P_1, m_J = 0\rangle$ state.

We use as an initial condition a thermal distribution of motional states with mean vibrational level $\bar{n} = 1$. We set $\Omega = \Gamma/500$, and perform the simulation for 1000 Rabi cycles. This ensures that we reach the system's steady state, and that the effect of the finite Ω on the final motional state populations is negligible.

We note that in this work, we do not perform a simulation of cooling in the three directions of space, even though we take into account the dipole emission pattern in 3D. We assume

that there is no thermalization between the different axes; if such cross-axis thermalization did exist it would lead to a higher temperature than we find here.

Idealized sideband cooling

We consider sideband cooling on an idealized cooling transition, which is defined as having an infinitely narrow linewidth ($\Gamma = 0$), but still having the ability to decay to the spin ground state. We set $\Omega = 2\pi \times 2.5$ kHz and assume that the difference in energy between the spin states is given by a wavelength of 698 nm, which would correspond to cooling on the 3P_0 state of strontium. We initialize the system in a thermal distribution with $\bar{n} = 1$, and perform the following simulation. We evolve the system's state for a duration corresponding to a sideband π pulse. We then project the state onto $|\downarrow\rangle$ in order to mimic spontaneous decay. During this event, a momentum kick is provided to the atom following the radiation dipole pattern described earlier. We repeat this process 100 times, which ensures that we reach the system's steady state.

Idealized erasure correction cooling using replacement

The idealized erasure correction cooling simulations are similar to the idealized sideband cooling. We set $\Omega = 2\pi \times 2.5$ kHz and initialize the system in a thermal distribution with $\bar{n} = 1$. We then perform the following simulation. Starting in $|\uparrow\rangle$, we evolve the system's state for a duration corresponding to a sideband π pulse. In order to simulate a perfect erasure detection and replacement, we get rid of the population in $|\downarrow\rangle$, and re-normalize the population of the various motional levels in $|\uparrow\rangle$. We repeat this process 100 times, which ensures that we reach the system's steady state.

Noisy simulation of erasure correction cooling using replacement

We finally describe the simulation of erasure correction cooling using replacement including relevant imperfections. We solve the Master equation described earlier in which we take into account the finite lifetime of the 3P_0 state, and the tweezer-induced heating rate, which we implement by adding a jump operator acting on the motional states; the tweezer-induced heating is independently measured as an increase in mean vibrational number by $3.0(4) \text{ s}^{-1}$. We initialize the system in a thermal distribution with $P'_0 = 0.77$, and simulate the sequence described in Fig. 3.13. We assume a perfect imaging fidelity in the erasure detection but take into account the tweezer-induced heating during the imaging and replacement times.

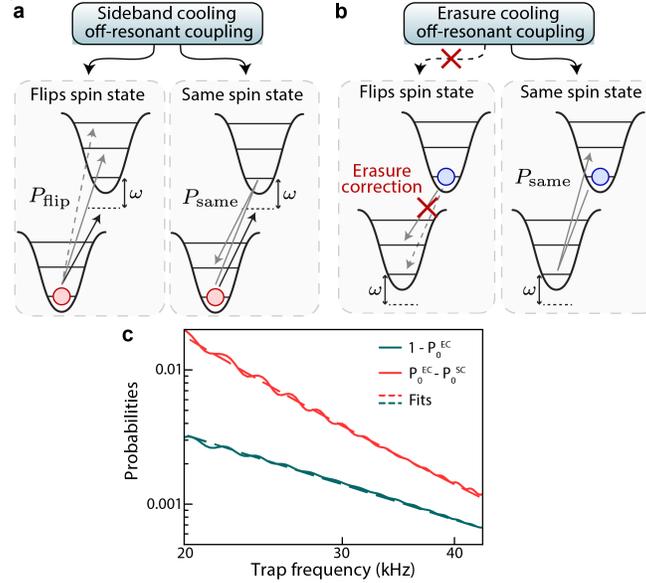


Figure 3.15: Limits of sideband cooling and erasure correction cooling **a**, The limit in sideband cooling efficiency is due to the off-resonant coupling to motional levels which are not the ground state ($n > 0$), either (i) in the opposite pseudo-spin state as the initial one $|n, \uparrow\rangle$, or (ii) in the same spin state $|n, \downarrow\rangle$ via an off-resonant two-photon process. **b**, In the case of erasure correction cooling, off-resonant coupling to the opposite spin state $|n, \downarrow\rangle$ is detected and hence does not contribute to the finite efficiency. The limit is the off-resonant two-photon transition to $|n, \uparrow\rangle$. **c**, Simulated efficiency of ideal ECC $1 - P_0^{\text{EC}}$ (red line), and difference between sideband cooling and ECC efficiencies $P_0^{\text{EC}} - P_0^{\text{SC}}$ (green line), as a function of trap frequency ω . We fit the data using $f(\omega) = A(1/\omega)^B$, with A and B as free parameters (dashed lines). The obtained exponents are $B = 1.9(1)$ for $1 - P_0^{\text{EC}}$ and $B = 2.9(1)$ for $P_0^{\text{EC}} - P_0^{\text{SC}}$, which is consistent with our expectations.

3.5.4 Fundamental limits simulations

Previously, we experimentally showed that erasure correction cooling using the replacement method leads to higher P_0 than a simulation of idealized sideband cooling (SC), for a specific range of parameters. Here we explain why, on a fundamental level, this new cooling mechanism leads to a larger P_0 than SC for any trapping frequency. We further determine the scaling of the ECC efficiency in the sideband resolved and Lamb-Dicke regimes, as a function of the trap frequency.

We start by describing the fundamental limits of sideband cooling. We assume that the atoms are initialized in $|0, \downarrow\rangle$, and that we drive a red sideband transition in order to cool the atoms. In this situation, sideband cooling has a finite probability to actually heat the atoms via off-resonant couplings to any motional state $|n, \uparrow\rangle$. Note that in the deep Lamb-Dicke regime, spontaneous emission does not lead to a change in motional state, and thus off-resonant coupling to $|0, \downarrow\rangle$ does not impact the finite cooling efficiency.

These off-resonant couplings affect the system in two qualitatively different ways (Fig. 3.15A): (i) the off-resonant coupling flips the initial pseudo-spin state to reach $|n, \uparrow\rangle$ with probability P_{flip} , and (ii) an off-resonant two-photon process which transfers the atoms in $|n, \downarrow\rangle$ and $n > 0$, with probability P_{same} . The values of P_{flip} and P_{same} depend on the Rabi frequency, Lamb-Dicke factor, and trap frequency, and are detailed below. The eventual fundamental limit in SC is therefore the sum of these two contributions:

$$P_0^{\text{SC}} = 1 - P_{\text{flip}} - P_{\text{same}},$$

where P_0^{SC} is the fundamental limit of the $|0\rangle$ population with sideband cooling.

We now turn to the case of ECC, and consider off-resonant coupling out of the initial state $|0, \uparrow\rangle$ (see Fig. 3.15B). For ECC, off-resonant coupling to opposite spin states $|n, \downarrow\rangle$ does not impact the ECC cooling efficiency: any atom transferred to $|\downarrow\rangle$ is detected and replaced, which can be effectively seen as $P_{\text{flip}} = 0$. The fundamental limit is thus given only by the pseudo-spin conserving, off-resonant two-photon transition $|0, \uparrow\rangle \rightarrow |n, \uparrow\rangle$, which cannot be converted into an erasure in our scheme. We therefore get:

$$P_0^{\text{EC}} = 1 - P_{\text{same}},$$

where P_0^{EC} is the fundamental limit in $|0\rangle$ population with erasure correction cooling.

We note that the value of P_{same} is the same in both cooling methods, as the action employed to cool the atoms – a red sideband pulse – is exactly the same for both methods. We therefore obtain that

$$P_0^{\text{EC}} - P_0^{\text{SC}} = P_{\text{flip}},$$

meaning that P_0^{EC} is an upper bound for P_0^{SC} .

Scaling of P_{flip} and P_{same} with ω

We further comment on the scaling of P_{flip} and P_{same} with the trap frequency ω , in the sideband resolved ($\Omega \ll \omega$) and Lamb-Dicke ($\eta(2\bar{n} + 1) \ll 1$) regimes. For simplicity, we will assume in the following that the atoms are initialized in the $|\downarrow\rangle$ state and that we drive a red sideband transition to the $|\uparrow\rangle$ state. Note that the results obtained below are also true if the atoms are initialized in $|\uparrow\rangle$.

Under these assumptions, P_{flip} to lowest order arises from off-resonant coupling via the $|0, \downarrow\rangle \rightarrow |1, \uparrow\rangle$ transition, with strength $\eta\Omega$ and detuning proportional to ω . We therefore obtain

$$P_{\text{flip}} \propto \frac{(\eta\Omega)^2}{\omega^2} \propto \frac{\Omega^2}{\omega^3},$$

where we used $\eta \propto 1/\sqrt{\omega}$.

P_{same} to lowest order arises from off-resonant coupling via the $|0, \downarrow\rangle \rightarrow |0, \uparrow\rangle$ transition, with strength $(1 - \eta^2)\Omega \simeq \Omega$ and detuning proportional to ω . Once this off-resonant transition happens, the on-resonance red sideband transition $|0, \uparrow\rangle \rightarrow |1, \downarrow\rangle$ transfers the atoms back to the same spin state, but in a higher motional level. We therefore obtain

$$P_{\text{same}} \propto \frac{\Omega^2}{\omega^2}.$$

We can therefore extract the scaling of idealized sideband cooling and erasure correction cooling in the sideband resolved and deep Lamb-Dicke regimes. In particular, we obtain

$$1 - P_0^{\text{EC}} \propto \frac{\Omega^2}{\omega^2},$$

and

$$P_0^{\text{EC}} - P_0^{\text{SC}} \propto \frac{\Omega^2}{\omega^3}.$$

We verify these scalings using our numerical simulations results of idealized SC and ECC (Fig. 3.15v), which are also presented in Fig. 3.14c. We consider the ECC deficiency $1 - P_0^{\text{EC}}$ (green solid line) as well as the difference between SC and ECC efficiencies $P_0^{\text{EC}} - P_0^{\text{SC}}$ (red solid line). In both cases, we fit our results with the function $f(\omega) = A(1/\omega)^B$ (dashed lines), where A and B are free parameters. The obtained exponents are $B = 1.9(1)$ for $1 - P_0^{\text{EC}}$, and $B = 2.9(1)$ for $P_0^{\text{EC}} - P_0^{\text{SC}}$. These results are consistent with the formulas derived above.

3.5.5 Erasure correction cooling on other platforms

Here we show that ECC is, in principle, applicable to any tweezer platform, and detail which states can be used for the $|\downarrow\rangle$ and $|\uparrow\rangle$ when implementing ECC. Our present work uses strontium, and a generalization to other alkaline-earth(-like) atoms is trivial, so we focus on platforms using molecules and alkalis (Fig. 3.16).

As discussed previously, the two key ingredients to implement ECC are (i) driving sideband transitions between $|\downarrow\rangle$ and $|\uparrow\rangle$, and (ii) the ability to selectively detect $|\downarrow\rangle$ while $|\uparrow\rangle$ is shelved. These ingredients are present in both molecules and alkalis.

For molecules (Fig. 3.16a), good candidates are two different rotational states, such as $|\downarrow\rangle = X^2\Sigma^+, N = 1$ and $|\uparrow\rangle = X^2\Sigma^+, N = 3$ in calcium fluoride. Sideband transitions can be achieved through Raman processes, and selective imaging of $|\downarrow\rangle$ is achieved through driving the closed-transition to $A^2\Pi_{1/2}$, for which $|\uparrow\rangle$ is shelved.

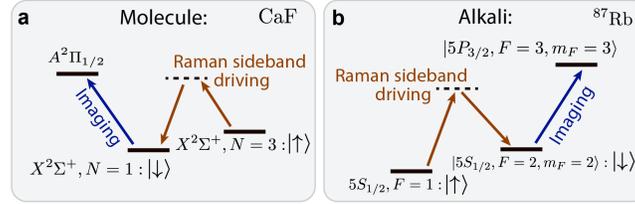


Figure 3.16: **Erasure correction cooling in other platforms.** **a**, Good candidates in molecules, using the example of calcium-fluoride. **b**, Good candidates in alkalis, using the example of rubidium.

For alkalis (Fig. 3.16b), good candidates are two different hyperfine states, such as $|\downarrow\rangle = |5S_{1/2}, F = 2, m_F = 2\rangle$ and $|\uparrow\rangle = |5S_{1/2}, F = 1, m_F = 1\rangle$ in rubidium 87. Sideband transitions can be achieved through a Raman process, and selective imaging of $|\downarrow\rangle$ is achieved through driving the closed-transition to $|5P_{3/2}, F = 3, m_F = 3\rangle$, for which $|\uparrow\rangle$ is shelved.

In principle ECC is also applicable to ion-based platforms, for which post-selection on detected erasures has been recently reported [170]. Furthermore, ECC should also work for neutral atoms and molecules trapped in optical lattices by using the reset method combined with local addressing in order to perform the local rotations [172].

3.5.6 Erasure cooling in the axial direction

The erasure cooling results presented previously concern one of the tightly-confining radial direction. We here show that the method also applies to the axial direction, enabling this new cooling mechanism to be used to perform 3D cooling in principle. The overall method is the same as described earlier, and the only difference lies in the direction of the sideband driving beam with respect to the tweezer axis. We use a beam that propagates in the same direction as the tweezer, allowing us to drive sideband transitions on the axial direction of the tweezers, for which the trapping frequency is $2\pi \times 10.7$ kHz.

The results are shown in Fig 3.17. We start with a motional ground state population $P'_0 = 0.40(2)$ after Sisyphus cooling, which is lower than the radial direction due to the lower trapping frequency. Applying erasure correction cooling, we obtain $P_0 = 0.982(+9/-18)$ after erasure detection and post-selection, and $P_0 = 0.925(+13/-25)$ with active correction using the replacement method. These results are also better than what can be achieved using sideband cooling on the 3P_1 transition for this trapping frequency, showing a practical advantage in using erasure correction cooling over sideband cooling.

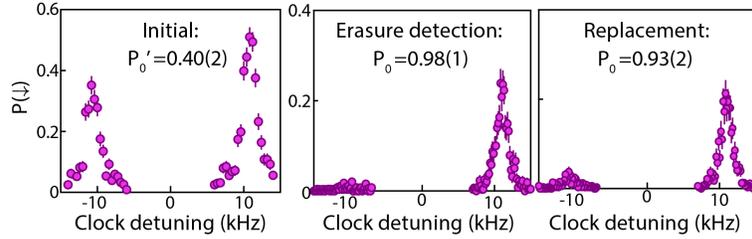


Figure 3.17: **Erasure correction cooling in the axial direction.** Probability P_0 for an atom to be in the $|0\rangle$ state measured via sideband spectroscopy, in the axial direction of the tweezers. We show the results without erasure correction cooling, performing erasure conversion with post-selection (not correction), and with correction using the replacement method.

3.6 DARK-STATE ENHANCED LOADING

Finally, we detail another application related to erasure conversion, namely using our ability to non-destructively locate defects in the array to perform feedback to improve array loading. This is a measurement- and feedback-based approach, and thus it is broadly reminiscent of classical erasure correction techniques like atomic rearrangement.

3.6.1 Overview

The central problem at hand is that while neutral atoms and molecules trapped in optical tweezers have become a powerful tool for quantum science, the maximum achievable system sizes of such arrays are often limited by the stochastic nature of loading into optical tweezers. This limits most experiments to a typical loading probability of only 50%. Here we present a species-agnostic method for dark-state enhanced loading (DSEL) based on real-time feedback, long-lived shelving states, and iterated array reloading. We demonstrate this technique with a 95-tweezer array of ^{88}Sr atoms, achieving a maximum loading probability of 84.02(4)% and a maximum array size of 91 atoms in one dimension. Our protocol is complementary to, and compatible with, existing schemes for enhanced loading based on direct control over light-assisted collisions [104, 173–177], and we predict it can enable close-to-unity filling for arrays of atoms or molecules.

In typical experiments, a magneto-optical trap (MOT) or large-waist optical dipole trap (ODT) containing a cloud of atoms is overlapped with a set of tweezer potentials which each capture an ensemble of atoms. Traps are then illuminated to induce light-assisted collisions and pairwise loss [35, 178, 179], resulting in a typical loading fidelity of 50% ($\sim 35\%$ for molecular tweezer arrays [180, 181]). Through direct control of the collisional process [173, 174], the array-averaged filling has been raised as high as 74-80% for arrays

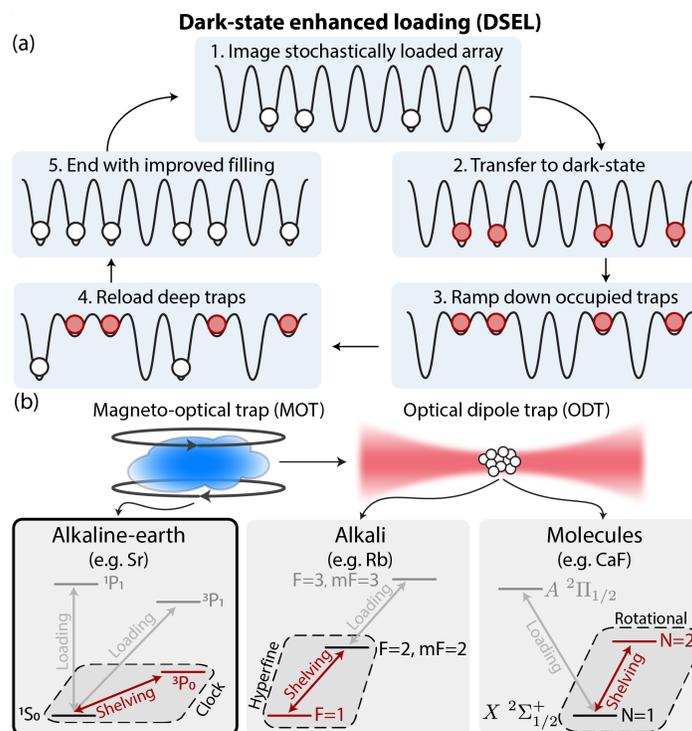


Figure 3.18: Dark-state enhanced loading a. A scheme for dark-state enhanced loading (DSEL) of an optical tweezer array. (1) Atoms (or molecules) are initially stochastically loaded into a tweezer array, and imaged to determine their locations. (2) Atoms are transferred to a long lived shelving state which is dark to the loading dynamics, and (3) occupied traps are ramped down in depth to limit reloading. (4) The loading procedure is repeated, stochastically reloading the unoccupied deep traps while leaving the shelved dark-state atoms unaffected, and (5) all atoms are returned to the ground state for an effective increase in filling fraction. The procedure can be repeated for successively higher filling fraction. **b.** DSEL is feasible to implement in arrays of alkaline-earth atoms (clock state shelving, e.g. Sr), alkali atoms (hyperfine state shelving, e.g. Rb), or molecules (rotational state shelving, e.g. CaF); the latter two both require tweezers be loaded from an ODT, rather than directly from a MOT. In this work we show an experimental demonstration with the alkaline-earth species ^{88}Sr .

of alkali species [175–177] and 93% for Yb [104]. While tweezer rearrangement can be used to eliminate remaining holes, the success probability and rearrangement time both scale poorly with the number of initial defects in the array [81, 89, 90], making methods for improving the initial loading not just impactful in terms of increasing atom number, but also in terms of increasing overall operational fidelity.

It is this context that we present our new, species-agnostic protocol for dark-state enhanced loading (DSEL), which we note is the first such demonstration of any enhanced loading protocol for strontium atoms. Our protocol is illustrated in Fig. 3.18a. In the initialization step, an array of atoms is loaded stochastically and imaged to identify the occupied traps.

Atoms are then shelved into a state which is both dark to the loading dynamics and long-lived. Occupied traps are ramped down in intensity through real-time feedback, and the process is restarted by regenerating the MOT or ODT which stochastically reloads the deep, unoccupied traps, while leaving the dark-state atoms in shallow traps unperturbed. The result is a theoretically 25% higher loading fraction (for an initial loading probability of 50%); the process can then be repeatedly infinitely for successively higher loading fidelity. We note that iterative loading techniques have been used to demonstrate continuous reloading of dual-atom tweezer arrays, though without increasing the filling fraction of either species [140]. The ingredients necessary to implement DSEL are accessible in arrays of alkaline-earth atoms, alkali atoms, or molecules (Fig. 3.18b), as we discuss further below.

3.6.2 Method

The experimental setup we use to implement DSEL is our standard array of ^{88}Sr atoms [69, 85, 99]. In brief, atoms are initially captured by a 461 nm (blue) MOT based on the $^1S_0 \leftrightarrow ^1P_1$ transition (Fig. 3.19a). With a small probability, some atoms can decay from $^1P_1 \rightarrow ^1D_2$ [99] and then into either 3P_2 or 3P_1 [182]. Atoms in 3P_2 are repumped with 707 nm light through the 3S_1 state to 3P_1 ; from there they can decay back to 1S_0 , closing the MOT cycle. Atoms are typically also repumped out of 3P_0 as there is a pathway from $^3P_2 \rightarrow ^3P_0$ during repumping.

Atoms are transferred from the blue MOT to a cold and dense 689 nm (red) MOT based on the narrow $^1S_0 \leftrightarrow ^3P_1$ transition. From there, an acousto-optic deflector (AOD) generates an arbitrary one-dimensional configuration of optical tweezer traps (813 nm); by tuning the electrical input to the AOD we control the number of tweezers, their spacing, and their relative intensities [81]. With a typical depth of 470 μK , optical tweezers are overlapped with the red MOT, loaded with an ensemble of atoms, and then illuminated with 689 nm light to induce light-assisted collisions [183, 184]. The transitions to the 1P_1 and 3P_1 states can then be respectively used to image and cool the atoms with high fidelity [85].

A key feature of the low-lying energy landscape of Sr is the 3P_0 state. This state is metastable [108], enabling applications in metrology [69, 110, 111, 185, 186] and it underlies proposals for alkaline-earth based neutral atom quantum computers [68, 187]. When in 3P_0 , atoms are dark to light involved in imaging, cooling, and MOT loading. In (Fig. 3.19b) we demonstrate this concept directly with three consecutive single-shot fluorescence images. In the first, all atoms are in the absolute ground state, and fluoresce under the 461 nm imaging light; during all images, 689 and 707 nm light is also active for cooling and repumping, respectively. Second, the central atom is locally excited to 3P_0 , making it dark to the

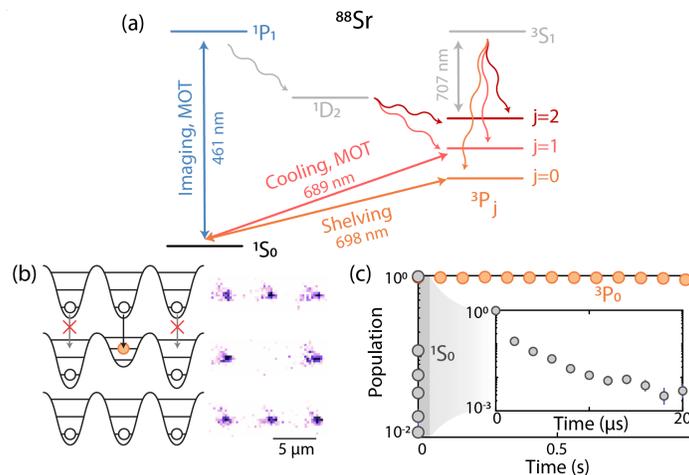


Figure 3.19: **DSEL in strontium.** **a.** Low-lying state manifold for ^{88}Sr . Of particular importance is the 3P_0 state which is long-lived, and dark to imaging, cooling, and MOT dynamics. **b.** Consecutive, single-shot fluorescence images of: (top) all atoms initially in the ground state; (middle) select atoms locally shelved via frequency-selective motional sideband driving to 3P_0 , and thus hidden from imaging; (bottom) all atoms brought back to the ground state, indicating the survival of the 3P_0 atoms. In all images 461, 689, and 707 nm light is illuminating the array. **c.** 3P_0 state population (orange) at 8 μK trap depth while the blue and red MOT light is on, showing a 23(4) s lifetime. In contrast, atoms in 1S_0 (grey) at this trap depth are quickly heated and expelled from the trap in less than 2 μs with the MOT light on (inset).

subsequent image. Local addressing is accomplished by lowering the central trap frequency by a factor of two and driving the motional blue sideband with a frequency selective pulse. Ramped down trap frequencies are ~ 40 kHz, relative to ~ 80 kHz for the non-ramped down traps, and a Rabi frequency of 2 kHz. Finally, all atoms are repumped to the ground state and imaged once more, demonstrating that the 3P_0 atoms have survived.

The lifetime of atoms in the 3P_0 state is limited by Raman scattering from the high-intensity trapping light [133]. We find a lifetime of 23(4) s for 8 μK deep traps while light used for generating the red and blue MOTs is active (Fig. 3.19c) and 0.58(4) s at our typical trap depth of 470 μK (not shown). At both trap depths the lifetime is unaffected by disabling the MOT light (not shown). In contrast, scattering from the MOT light quickly heats and ejects atoms in the 1S_0 state, limiting their trapped lifetime in 8 μK deep tweezers to less than 2 μs (Fig. 3.19c); this amounts to an effective increase in lifetime by a factor of $\sim 10^7$ when shelving into 3P_0 .

Isolation from other transitions, in combination with a long lifetime, makes 3P_0 an ideal candidate for the dark-state required by DSEL. However, during the standard MOT loading cycle, 3P_0 is explicitly repumped. In practice, we find that under our typical conditions,

loading without the 3P_0 repumper incurs only a 0.7(4)% drop in loading probability; through optimization of the blue MOT density it is likely that this minor penalty could be further reduced.

3.6.3 Results

We show the results of DSEL for a 95-tweezer array in Fig. 3.20a. With an initial loading probability of 45.88(5)%, after a single cycle of DSEL we achieve 68.43(5)%, which then saturates to 84.02(4)% after five cycles of DSEL (the zeroth cycle is the initial standard loading). We show exemplary fluorescence images from a single set of DSEL cycles in Fig. 3.20b, demonstrating the cycle-to-cycle growth of array filling, resulting in a maximum of 91 trapped atoms. Note that in Fig. 3.20b, atoms are additionally rearranged to the middle of the array between each cycle.

We see the effect of DSEL starts stagnating after around 2 cycles, limited by the probability that a shelved atom survives and still is present after a cycle of DSEL, which we designate as the cycle-to-cycle survival, p_0 . Assuming the loading process is only controlled by this cycle-to-cycle survival and by the per-cycle loading probability, f_0 , we predict the total loading probability, \bar{f} , after κ DSEL cycles is given by

$$\bar{f}(\kappa) = f_0 \sum_{j=0}^{\kappa} (p_0 - f_0)^j = f_0 \frac{1 - (p_0 - f_0)^{\kappa+1}}{1 - (p_0 - f_0)}, \quad (3.11)$$

where j is a cycle index. For the 95-tweezer array, we find our results are well described by Eq. (3.11) with $f_0 = 0.46$ and $p_0 = 0.93$ (Fig. 3.20a, inset).

For a smaller 39-tweezer array with improved atomic survival and dark transfer efficiency, the cycle-to-cycle probability is improved to $p_0 \sim 0.97$ (Fig. 3.20c, upper). We find p_0 is primarily controlled by imperfect transfer to 3P_0 (transfer to 3P_0 is performed coherently, as incoherent pumping leads to atomic heating, and increased loss.), decay of the atoms out of 3P_0 , and atom survival (Fig. 3.20c, upper). This is because if atoms are not transferred to 3P_0 , or if they decay back to the ground state, then they will be ejected from the traps during the MOT dynamics, leading to loss (Fig. 3.19c).

Traps are ramped down immediately following transfer in order to extend the lifetime of atoms in 3P_0 and limit this effect. However, traps cannot be ramped down arbitrarily low as the probability for atoms to remain trapped, regardless of electronic state, decreases below a certain trap depth (Fig. 3.20c, upper) due to their temperature [188]. Ultimately we find an optimum of these effects at a trap ramp down factor of 1/60 (a trap depth of 8 μK); at this

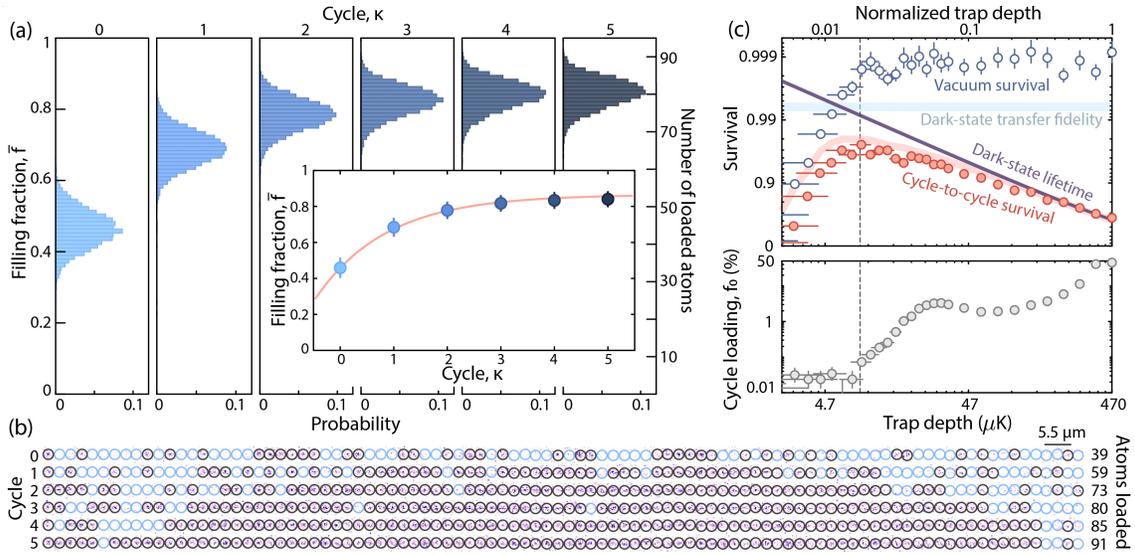


Figure 3.20: Continuous and record enhanced loading with DSEL. **a.** DSEL results in a 95-tweezer array. Shown in the main plot are histograms of the filling fraction of the array after each successive repeated loading cycle. The zeroth cycle is the standard, non-enhanced loading with a loading fraction of 45.88(5)%, while after five DSEL cycles a maximum of 84.02(4)% is achieved. Inset: Average and standard deviation of the filling fraction as a function of loading cycle (markers), compared against the prediction from Eq. (3.11) based on a per-cycle loading fidelity of 0.46, and a cycle-to-cycle survival probability of 0.93 (pink line, see text). **b.** Exemplary single-shot fluorescence images from a single set of DSEL cycles, showing the largest-achieved 91 atom array. Sites above the atom-detection threshold are circled in black, while those below the threshold are circled in blue. Note that atoms are additionally rearranged to the middle of the array between each cycle, though this is not required. **c.** Cycle-to-cycle survival (red markers) is limited by imperfect dark-state transfer fidelity (blue fill), dark-state lifetime (purple line), and vacuum-limited survival (open markers) during the reloading cycle, here 280 ms; x-error bars indicate uncertainty in the exact ramp-down trap depth due to a fluctuating bias on a photodiode used for stabilization. The combination of these effects is shown as a pink line. An optimum in survival emerges at a ramp-down factor of 1/60th (a trap depth of 8 μK). At this trap depth, the probability of reloading an already occupied tweezer is $\lesssim 0.1\%$ (bottom). Data in **c** are taken in a 39-tweezer array with higher fidelity survival and dark-state transfer compared to the 95-tweezer array in **a** and **b**.

trap depth we additionally find the probability of tweezer reloading is $\lesssim 0.1\%$ (Fig. 3.20c, lower), limiting the probability for new atoms to be loaded into already occupied traps.

Having demonstrated DSEL experimentally, in Fig. 3.21a we plot the results of Eq. (3.11) as a function of total number of cycles for various p_0 at a fixed $f_0 = 0.5$, and include the infinite cycle limit of:

$$\bar{f}(\infty) = f_0 \frac{1}{1 - (p_0 - f_0)}. \quad (3.12)$$

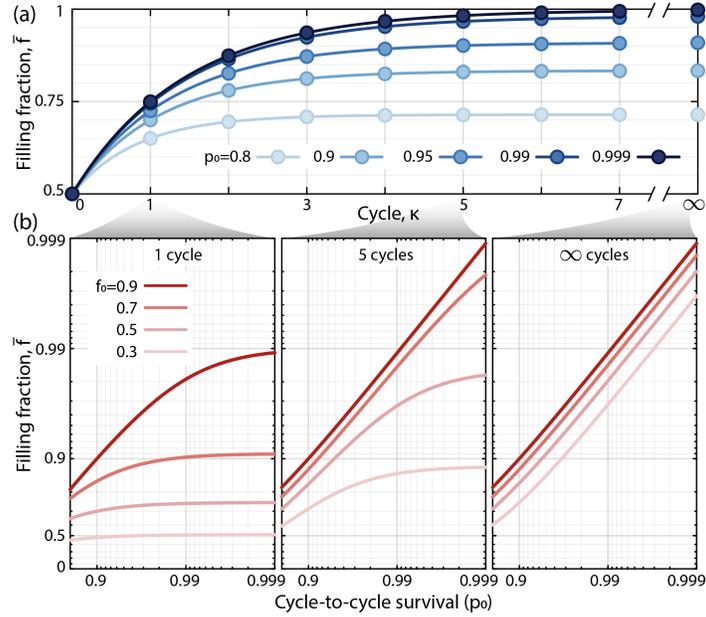


Figure 3.21: **Limits of DSEL.** **a** Predicted total filling fraction from Eq. (3.11) with varying number of DSEL cycles, κ , and a fixed per-cycle filling fraction, $f_0 = 0.5$. **b** Total filling fraction after DSEL for various cycle numbers, per-cycle filling fractions, f_0 , and cycle-to-cycle survival probability, p_0 ; f_0 values greater than 0.5 are achievable by combining DSEL with existing methods for directly controlling the collisional processes during loading [104, 173–177].

We find that the final filling fraction saturates, with the saturation level varying with p_0 . To understand this behavior, we plot the filling as a function of p_0 for 1, 5, and infinite cycles in Fig. 3.21b. There we show the results for various f_0 ranging from 0.3 to 0.9, given that f_0 can be increased with existing methods for enhanced loading based on control over light-assisted collisions [104, 173–177].

In the many cycle limit, we find that DSEL is limited primarily by p_0 ; for p_0 approaching 1, Eq. (3.12) reduces at lowest order to $\bar{f}(\infty) \approx 1 + (p_0 - 1)/f_0$. Even with perfect initial loading ($f_0 = 1$), this further simplifies directly to $\bar{f}(\infty) \approx p_0$. On the other hand, if $p_0 = 1$ then $\bar{f}(\infty) = 1$, no matter the value of f_0 . Intuitively, this is because the gain from DSEL stems from accurately learning information from atomic imaging to reduce loading entropy. If that information is erroneous, i.e. if p_0 is low, then there can be relatively little gain (and if $p_0 < f_0$, applying DSEL can even lead to a loss of atoms). On the other hand, the gain is maximized when the information gained from imaging the array is most accurate, i.e. when p_0 is maximized. With state-of-the-art imaging survival [85], dark-state transfer [34, 119], and optimized MOT loading times, $p_0 > 0.99$ is likely achievable.

3.6.4 Discussion

Though demonstrated here only for the alkaline-earth atom Sr, DSEL is feasible to implement for both alkali and molecular tweezer arrays (Fig. 3.18b). For the case of laser-coolable molecules, such as CaF, the MOT is typically transferred to an ODT to increase molecule densities [180]; molecules are then cooled into tweezer potentials via Λ -cooling on the $|X \ ^2\Sigma_{1/2}^+, N=1\rangle \leftrightarrow |A \ ^2\Pi_{1/2}, J=1/2, +\rangle$ transition. In this case, the dark-state for DSEL could be another long-lived rotational state, such as $N=2$, which is far-off resonant from this cooling process¹.

In the alkali case, such as with Rb, the dark-states would be a long-lived hyperfine manifold such as $F=1$; however, because it is difficult to form the initial MOT without repumping from these states, it may be necessary to first transfer from the MOT into an ODT, similar to the molecule case. Atoms can be cooled from the ODT into tweezers using e.g. the $|F = 2, m_F = 2\rangle \leftrightarrow |F = 3, m_F = 3\rangle$ transition, to which $F=1$ shelved atoms would be dark [189].

Loading directly from an ODT may prove beneficial in general, besides the particular application to the molecular and alkali cases. This is because once the ODT is formed, it can be moved away while the tweezer array is imaged, and then the two can be overlapped once more to perform DSEL without the need to repeat the entire MOT sequence. This could greatly reduce one of the main limitations of DSEL, namely the increased experimental runtime.

Importantly, DSEL can be implemented with practically no impact on experiment duty cycle by reusing atoms from one iteration of the experiment to the next. Typical experimental sequences end in a stage of readout, which can then be followed by transferring any remaining atoms to the dark-state before preparing the next experimental repetition. This effectively passes atoms from one experimental iteration to the next by implementing a 1-cycle version of DSEL. Such an approach has since been realized in Ref. [141], which essentially implements DSEL in a two-dimensional array of strontium atoms for thousands of cycles.

If a species has internal structure to its dark-state, coherence could feasibly be maintained within that manifold during DSEL, similar to an approach utilizing dual-species arrays [140]. For instance, the multiple hyperfine sublevels of 3P_0 in fermionic alkaline-earth species

¹We especially acknowledge discussions with molecule-experts Arian Jadbabaie, Loïc Anderegg, and Yicheng Bao, and thank them for their insights and suggestions

could be used to shelve atoms while performing mid-circuit replacement of lost atoms [68, 190].

We emphasize that though here we have implemented DSEL in one dimension, the basic principle is applicable to arbitrary-dimensional systems, as long as there is sufficient control over the trap depths of individual traps (or regions of traps) during real-time feedback. This is possible with either crossed AOD, time-multiplexed, or holographic techniques, as are typically used for two-dimensional array generation [99, 103, 191–194]. It can also be achieved with bichromatic trapping potentials [141].

We finally note a natural extension to DSEL by adding more tweezers to the array on each round of loading. Tweezer arrays are typically limited in scale by available laser power needed to perform loading and imaging with high fidelity, but following DSEL atoms can be held in shallow traps, allowing for the total laser power to be redistributed to generate more total traps, and thus increasing the array size. With such a protocol, effective loading fidelities greater than 1 (relative to the original number of traps) could be realized, and the requirement that $p_0 < f_0$ could be softened.

In summary, we have proposed and demonstrated a protocol for enhancing the filling fraction of an optical tweezer array through dark-state enhanced loading. This protocol achieved a filling fraction of 84.02(4)% in a 95-tweezer array and yielded an array with 91 atoms; to our knowledge this is the first such enhanced loading demonstrated with a Sr atom array, and the largest one-dimensional array of atoms in optical tweezers to date. Importantly, our protocol is scalable, complementary to existing enhanced loading schemes, and species-agnostic. Ultimately, consistent creation of larger arrays will have benefits for metrology, quantum simulation, and tests of quantum advantage in optical tweezer systems.

CHAPTER 4

Tools of a universal quantum processor

Having described the Rydberg Hamiltonian, its error sources, and means to mitigate those errors, much of the rest of this thesis will focus on what we can learn from such a system operating in an analog context. However, importantly, Rydberg atom arrays also can be operated in a *digital* modus, where continuous operations are replaced by a discrete, gate-based approach. While such digital quantum computers are currently limited by noise and compilation times, with the advent of error correction such systems will be universal in their capabilities. Even in the modern day, Rydberg atom arrays operating as digital quantum processors have shown impressive results [37–40] thanks to a slew of utilities such as high-fidelity single- and two-qubit gate operations [39, 74, 75, 95, 104, 154, 195], mid-circuit measurement [32, 40, 106, 139, 196, 197], and more.

In this chapter we will demonstrate these tools of universal quantum computation. In almost all prior cases with atoms, such demonstrations were performed with a long-lived hyperfine or nuclear qubit, while in our case our long-lived qubit will be composed of the absolute ground state and the metastable clock state. This is thus an *optical qubit*, which brings with it a distinct set of challenges and advantages. This is especially important because such optical transitions, especially in strontium [110, 111], constitute the basis for state-of-the-art optical atomic clocks [198]; by combining a digital quantum processor with such a pristine clock system, even greater metrological capabilities could be achieved [28, 39, 136, 147, 199–201].

Specifically, we first demonstrate arbitrary, parallel, single-site addressing for an optical transition via sub-wavelength controlled moves of tweezer-trapped atoms [28]. Our method is robust, and easily scalable to larger arrays while maintaining high-fidelities of ~ 0.999 , and with essentially zero crosstalk. Next we will demonstrate mid-circuit measurement

and state shelving by coherently manipulating the motional degrees of freedom in optical tweezers [32]. Our method is applicable to all tweezer platforms, but especially important for bosonic strontium arrays due to the lack of otherwise readily available shelving states. Then, we will demonstrate two-qubit gates on pairs of optical qubits via transient excitation to Rydberg states, and will show the formation of optical Bell pairs. We will combine these two-qubit gates with our motional shelving techniques to generate *hyper-entanglement* [202, 203] by preparing a simultaneous Bell state of motional and optical qubits. This is the first such demonstration of hyper-entanglement for any matter-based (i.e. non-photon) qubit, and could be the basis for a number of important tasks in quantum information science.

Finally, we combine this diverse and universal set of tools to demonstrate applications in quantum metrology for potentially exponentially improving the capabilities of atomic clocks [28]. First, we use our arbitrary local addressing scheme to implement single-shot, dual-quadrature readout of Ramsey interferometry using two atomic ensembles simultaneously, and show an enhancement of the usable interrogation time at a given error probability. This scheme can be even further improved through the use of many more ensembles of atoms evolving with different phase sensitivities, as has been long studied theoretically. Realizing such schemes could be done in many ways, including local dynamical decoupling [200, 201], mid-circuit measurement and reset [200, 201], and via creation of registers of entangled states [147, 199]. We first show the local dynamical decoupling technique, using local single-site control to evolve three ensembles with variable phase sensitivities. Then we discuss our ongoing work to combine this technique with entangled states of clock qubits, and show a proof-of-principle [39].

Ultimately our results demonstrate that digital quantum processing is not only viable in strontium-based atom arrays, but is practical and immediately useful. We foresee that in future years as this digital processing modus becomes more refined, state-of-the-art quantum metrology may be performed using versions of our techniques demonstrated here.

4.1 LOCAL ADDRESSING WITH ATOM SHIFTS

Here we show local control of optical transitions in a tweezer array clock [69, 112, 186] by using rearrangement techniques [37, 81, 89, 204–206] on atoms in superposition states to precisely control the position-dependent phase imprinted by light-matter interaction. The scheme [187, 207] is experimentally simple and highly robust as it relies solely on the relative stability of tweezer positions and does not involve any auxiliary addressing beams. Using this technique, we demonstrate arbitrary, parallel, single-site-resolved optical qubit rotations with high fidelity. Later in this thesis we will discuss this approach in the context of

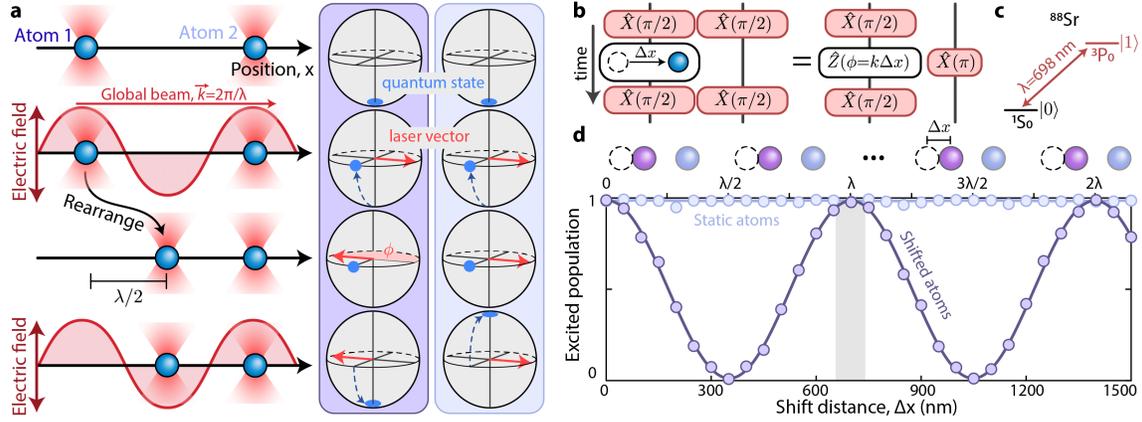


Figure 4.1: **Single-site addressing with movement-induced phase shifts.** **c.** We consider two atoms individually trapped in optical tweezers, both initially in the electronic ground state. Traveling light emitted from a global laser beam applies a $\pi/2$ rotation to both atoms, and is disabled, but remains phase coherent with the atomic transition. One of the atoms is then moved by half the laser wavelength, λ , from its initial position, rotating the effective local laser frame by an angle $\phi = \pi$. When the laser drive is restarted to apply another $\pi/2$ pulse, the moved atom now rotates back to the ground state, while the static atom rotates to the excited state. **c.** Control over the atom displacement, Δx , is equivalent to arbitrary local rotations of the laser drive by $\phi = k\Delta x$ about the \hat{Z} -axis. **c.** We implement this protocol with an array of ^{88}Sr atoms utilizing the ultra-narrow $^1S_0 \leftrightarrow ^3P_0$ transition with $\lambda = 698.4$ nm for global driving. **d.** Top: with an array of 39 tweezers in one dimension, we apply the protocol in **b**, shifting every odd site (purple markers) in the array while leaving all even sites static (blue markers) during the dynamics. Bottom: a sinusoidal oscillation emerges in the excited state population of the shifted sites, with a period of $699(1)$ nm.

building better atomic clocks, but our technique for implementing local, parallel rotations about arbitrary axes might also find use in neutral atom quantum computing platforms utilizing optical transitions [68, 187], where local coherent control of optical qubits has not been demonstrated before this work. More generally, our results point to a future of fully programmable neutral atom optical clocks that incorporate features of quantum computers [39].

The basic principle of our scheme is illustrated in Fig. 4.1a. We consider two atoms both initially in the ground electronic state, $|0\rangle$, interacting with a global laser beam characterized by wavevector $k = 2\pi/\lambda$, and wavelength λ , propagating along the array axis. With the globally applied laser, we create an equal superposition state of $|0\rangle$ and the excited state, $|1\rangle$; in a Bloch sphere picture, this corresponds to a $\pi/2$ rotation around the x -axis ($\hat{X}(\pi/2)$). The laser beam is then extinguished with an optical modulator, but remains phase coherent with the atomic transition. Using atom rearrangement techniques [37, 81, 89, 204–206], one of the atoms is shifted from its original position by Δx , applying an effective phase shift

of $\phi = k\Delta x$ (see Eq. (4.1)). In Fig. 4.1a, we first consider the special case of $\Delta x = \lambda/2$, or equivalently a π rotation around the z -axis ($\hat{Z}(\pi)$) for the shifted atom (Fig. 4.1a). Subsequently, we apply a second global $\hat{X}(\pi/2)$ rotation with the same laser as before; the shifted atom now rotates back to $|0\rangle$ because of the movement-induced phase shift, while the unmoved atom completes its rotation to $|1\rangle$.

The main principle behind this scheme is a locally controlled change of the relative phase between the atomic dipole-oscillation and the phase of the laser while the atom is in a superposition state; in essence, our scheme realizes a locally controlled Ramsey sequence with global driving. Similar techniques have been used in the context of ion trap experiments with two ions [207], but not in a scalable fashion, as is possible with tweezer arrays [187].

We show an experimental demonstration with our ^{88}Sr optical tweezer array experiment [29, 69, 99]. We employ a one-dimensional array of 39 optical tweezers generated via an acousto-optic deflector (AOD) driven by an arbitrary waveform generator (AWG). This allows for precise control over the relative tweezer positions at the nanometer level, enabling arbitrary $\hat{Z}(\phi)$ rotations (Fig. 4.1b). Global driving is performed on the ultra-narrow $^1S_0 \leftrightarrow ^3P_0$ optical clock transition with a transition wavelength of $\lambda = 698.4$ nm (Fig. 4.1c).

In a first experiment, we apply an $\hat{X}(\pi/2)$ operation globally to the entire array, then shift every odd site by the same distance, Δx , apply another global $\hat{X}(\pi/2)$ rotation, and finally measure the excited state population in both the shifted and unshifted sub-arrays. The excited state population of shifted atoms, P_s , shows sinusoidal oscillations with a period of 699(1) nm as a function of Δx , consistent with $\phi/2\pi = \Delta x/\lambda$, where λ is the transition wavelength. The quoted error on this measurement is purely statistical, and ignores potential systematic error arising from the independent distance calibration performed with an optical resolution test target. We note that the present measurement is likely a far more precise and accurate distance calibration tool, and could find use as an effective *in-situ* laser-based ruler with applications for precision determination of distance-dependent inter-atom effects, such as Rydberg interactions.

To quantify the phase shift fidelity, we focus on a narrow region around $\Delta x = \lambda$, corresponding to a $\hat{Z}(2\pi)$ rotation (Fig. 4.2a). A quadratic fit to $P_s(\Delta x)$ shows a maximum value of $P_s = 0.9940(5)$ (not corrected for state preparation and measurement (SPAM) errors), consistent with the mean excited state population of unshifted atoms, $P_u = 0.9942(2)$, in the same range. Correcting for the bare $\hat{X}(\pi)$ fidelity (red dashed line) of 0.9956(1) shows the shift operation is performed with a fidelity of 0.9984(5). We note that applying SPAM correction on the bare fidelities maintains the shift fidelity largely unchanged as it is calculated from the ratio of the two. The ratio of the shifted to unshifted fidelities is 0.9998(5),

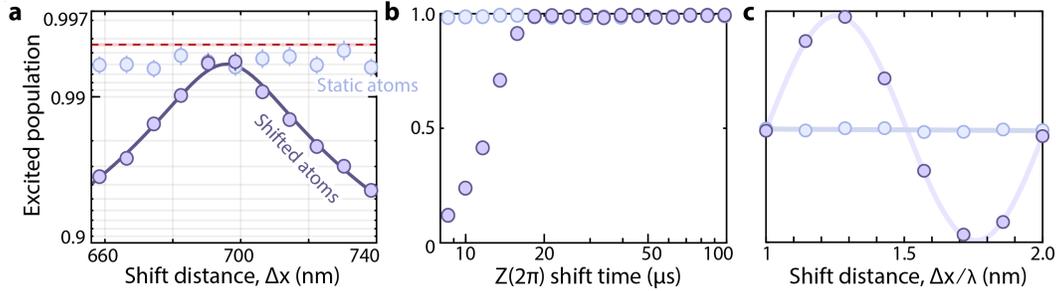


Figure 4.2: **Fidelity of local addressing c.** Focusing on the region around $\Delta x = \lambda$ (grey shaded region in Fig. 4.1d), we find the shifted atom shows no measurable loss in fidelity compared to the unshifted atoms. Correcting for the bare fidelity for performing a global $\hat{X}(\pi)$ rotation (red dashed line, 0.9956(1)), we find the shift operation is performed with a fidelity of 0.9984(5). The ratio of the shifted to unshifted fidelities is 0.9998(5), suggesting that the dominant source of error comes from global laser phase noise during the finite wait time required to perform the shift, rather than the movement itself. **c.** The shift to apply a $\hat{Z}(2\pi)$ rotation can be performed without noticeable loss of fidelity down to shift times of $\sim 20 \mu\text{s}$; data in **e** are taken with a shift time of $32 \mu\text{s}$, in addition to an extra wait time of $34 \mu\text{s}$ to account for finite jitter in the control timings. **c.** From data in Fig. 4.1d, we find the crosstalk to the static atoms is 0.1(2)%, consistent with 0. To verify this, we check for any residual linear phase shifts by repeating that experiment with a global phase shift of $\pi/2$ on the final Ramsey pulse which yields linear sensitivity to small phase shifts. We find the unshifted atom's phase to still be consistent with zero over a range of more than one wavelength.

suggesting that the dominant source of error comes from global laser phase noise during the finite wait time required to perform the shift, rather than the movement itself. We study the fidelity to perform the $\hat{Z}(2\pi)$ rotation as a function of the shift time (Fig. 4.2b), and find that the fidelity remains constant down to shift times of $t_s = 20 \mu\text{s}$; data in Fig. 4.2a were taken with $t_s = 32 \mu\text{s}$, plus an additional $34 \mu\text{s}$ of wait time to account for jitter in the subsequent control timings. Timing improvements could be made by implementing more sophisticated atom movement profiles [208]. Importantly, for all shift distances in Fig. 4.1d, the excited state population of the neighboring unshifted atoms is nearly constant, showing crosstalk of only 0.1(2)% (Fig. 4.2c).

Arbitrary rotation patterns can be imprinted on the array by shifting all of the atoms by varying distances such that rotations about the z -axis with tweezer-resolved phase, ϕ_j , are applied (Fig. 4.3a). We show the results of time-resolved Ramsey spectroscopy for four different choices of single-site addressing patterns, demonstrating arbitrary, site-revolved, and parallel \hat{Z} rotations (Fig. 4.3b). Such addressing patterns could be used to negate variations in the transition frequency across the array, for instance due to gradients in magnetic field or from the finite differences in tweezer wavelengths as generated by an AOD [69].

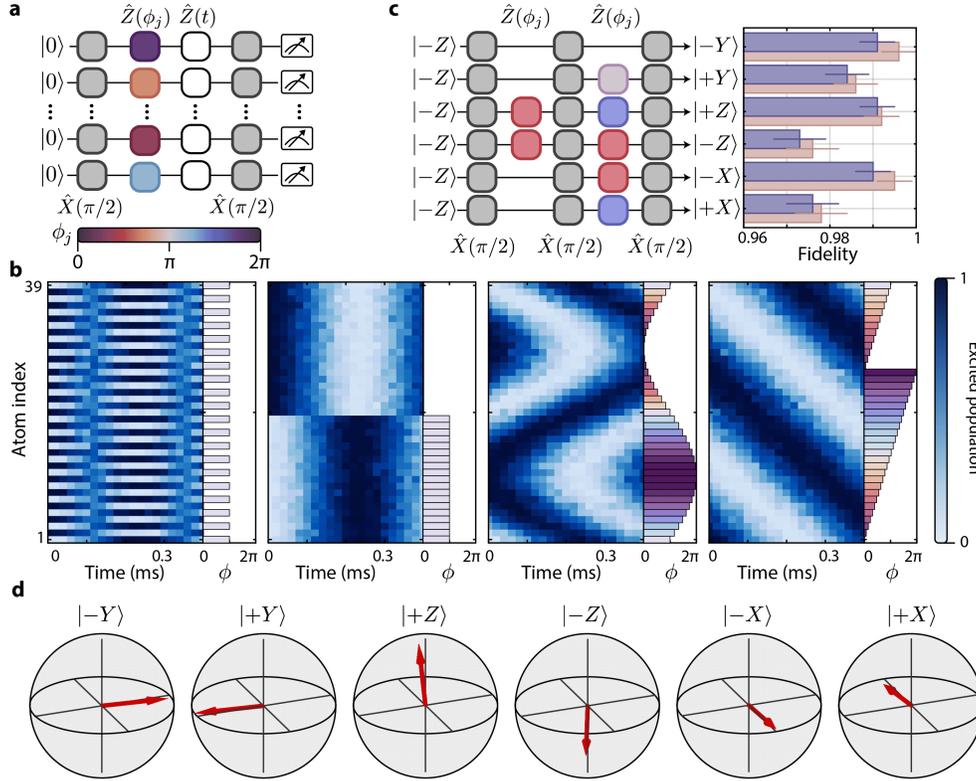


Figure 4.3: Arbitrary, parallel, local rotations. **c.** We implement site-resolved phase shifts, ϕ_j , during the dark time, t , of standard Ramsey interrogation by inserting arbitrary and parallel shifts of various distances to the array of atoms. **c.** Results of this operation as a function of Ramsey time (x -axis) for different tweezers in the array (y -axis). The corresponding programmed phase-shift pattern is shown on the right of each subfigure. **c.** By applying multiple global $\hat{X}(\pi/2)$ pulses (grey blocks), in tandem with local movement shifts (same color scale as in **a**), arbitrary local rotations can be performed. We show a demonstration by rotating an array of six atoms, initially in the $|0\rangle = |-Z\rangle$ state, in parallel to the six cardinal states ($|-Z\rangle, |+Z\rangle, |-Y\rangle, |+Y\rangle, |-X\rangle, |+X\rangle$), achieving an average fidelity of 0.984(2) (blue bars), and 0.987(2) after state preparation and measurement (SPAM) correction (tan bars), limited by global $\hat{X}(\pi/2)$ fidelity and decoherence during the time needed for movement. **d.** Bloch sphere visualizations of the states measured with quantum state tomography in **c**.

Combining these single-site $\hat{Z}(\phi_j)$ rotations with a series of global $\hat{X}(\pi/2)$ pulses allows for rotations about *any axes*, not just the z -axis. As a demonstration (Fig. 4.3c,d), we choose a set of 6 contiguous atoms, initially in the ground state (denoted here as $|-Z\rangle$), and rotate them each in parallel into the six states $|-Z\rangle, |+Z\rangle, |-Y\rangle, |+Y\rangle, |-X\rangle, |+X\rangle$, with an average fidelity of 0.984(2) (0.987(2) SPAM-corrected), as determined by state tomography (described below). The dominant limitations to this value are likely from global drive infidelity and dephasing during the finite shift times. We further note that this technique can readily be extended to realize robust pulses not just to realize the cardinal states [209]

We note that while here we have demonstrated our protocol on a one-photon optical transition, it could be used to induce a similar effect for two-photon Raman transitions, for instance between hyperfine states, assuming the two beams are counter-propagating. Further, the movement-induced phase-shifts employed here rely solely on a relative change in tweezer position, in contrast to alternative techniques that apply additional addressing beams [38, 154, 172, 210, 211], where the phase shift is proportional to a local addressing beam's intensity. While the addressing beam intensity and alignment are prone to drifts on experimental time scales, relative atom movements are ultimately derived from the radiofrequency electronic output of an AWG, which is precise, consistent, and robust. We emphasize our results did not utilize noise-compensating composite pulse sequences and that all data were taken without any system realignments or recalibrations of the atom movements.

4.1.1 Light-matter interaction Hamiltonian

After applying the rotating-wave-approximation, the light-matter interaction Hamiltonian considered in this work is given by [129]

$$\hat{H} = \frac{\hbar|\Omega|}{2}(|0\rangle\langle 1|e^{ikx} + |1\rangle\langle 0|e^{-ikx}), \quad (4.1)$$

where Ω is the Rabi frequency, k is the global laser wavevector, x is the atomic position along the beam propagation axis, $|0\rangle$ is the ground state and $|1\rangle$ is the excited state. Atom displacements by Δx correspond to phase shifts of $\phi = k\Delta x$, as described in the main text.

Note the choice of the phase for the initial $\hat{X}(\pi/2)$ rotation is a local gauge freedom, and thus all atoms in the array can be said to experience the rotation about the same local axis, e.g. the x -axis, despite the spacing between atoms generically not being perfectly commensurate with the driving wavelength. When the atom is shifted, it can be thought of as changing $\hat{X} \rightarrow \hat{X} \cos(\phi) + i\hat{Y} \sin(\phi)$. For the case of only a single global \hat{X} rotation after the movement, this is equivalent to an effective $\hat{Z}(\phi)$ rotation of the quantum state; however, in general if multiple global \hat{X} operations are performed, the equivalence with an effective \hat{Z} rotation breaks down.

4.1.2 Operation fidelities

Error modeling suggests that the global $X(\pi)$ fidelity of 0.9956(1) is primarily limited by: measurement errors (see below), finite temperature (the average motional occupation of the atoms along the radial axis is $\bar{n} \sim 0.2$, leading to a 2×10^{-3} infidelity), and frequency noise

on our laser; the latter of these we believe is also the dominant limitation to our Ramsey coherence time. The Rabi frequency is ~ 2.5 kHz for all measurements in this work, which allows for fast operations compared to the timescale of decay from 3P_0 (~ 550 ms [33]). Measurement errors are dominated by the vacuum-limited atom survival during imaging (0.9995(4)), the imaging fidelity for detecting the presence of an atom (0.9997(2)), and the likelihood of ejecting atoms in 1S_0 from the tweezers to perform state discrimination (0.9967(5)) before the final imaging. We note that not all of these measurement errors contribute equally, and their relative importance generically depends on the amount of excited state population in the measured state.

The wavelength of the oscillation, λ_{osc} in Fig. 4.1d is found from fitting the excited state population of shifted atoms with a sinusoid of the form $A \sin(2\pi\Delta x/\lambda_{\text{osc}}) + B$, from which we determine the quoted value of $\lambda_{\text{osc}} = 699(1)$ nm. Per our independent calibration of distances within the array, we expect there is an additional ~ 5 nm of systematic uncertainty on this measurement. Further systematic uncertainty could arise if the beam propagation is not perfectly coaxial with the array. The crosstalk fidelity of 0.1(2)% is found by fitting the unshifted atom excited state populations in Fig. 4.1d with a sinusoid of the same period as was determined for the shifted atoms; the quoted crosstalk is then the amplitude, A , of this sinusoid. We also check for any residual linear phase shifts by repeating this experiment with a global phase shift of $\pi/2$ on the final Ramsey pulse which yields linear sensitivity to small phase shifts (Fig. 4.2c). We find the unshifted atoms phase to still be consistent with zero over a range of more than one wavelength.

To determine the fidelity of arbitrary local rotations (Fig. 4.3c), we perform quantum state tomography by reading out the produced states in the x -, y - and z - bases by rotating the state with a global $\pi/2$ pulse of a given global laser phase as necessary. The fidelity is then estimated with $F = \langle \psi_{\text{target}} | \rho | \psi_{\text{target}} \rangle$, where ρ is the experimental state determined by quantum tomography. Due to the choice of the six arbitrary rotations, the fidelity estimation coincides with the population of the excited state or the ground state along the corresponding axis. For instance, the fidelity of $|+Y\rangle$ is determined by the population of the excited state when the prepared state is rotated and measured in the y -basis. In order to access the intrinsic infidelity induced by arbitrary local rotations, we extract SPAM errors and correct them. SPAM sources are dominated by the same detection infidelities as for the global $\hat{X}(\pi/2)$ fidelity (see above), and the finite SPAM-corrected readout $\pi/2$ pulse fidelity of 0.9982(4).

4.2 MID-CIRCUIT READOUT

Over the last few years, tremendous efforts have been pursued in a wide array of quantum platforms to develop efficient quantum information carriers [145, 212–215]. In particular, for neutral atom based platforms, quantum information has so far been encoded in electronic or nuclear states [36, 75, 95, 216]. However, there is another manifold in which information can be encoded: the *motional state* of a trapped atom. Here we demonstrate how mid-circuit measurement can be accomplished in arrays of strontium atoms via shelving coherent spin superposition states into the motional manifold where they are protected for tens or even hundreds of milliseconds.

The motion of atoms trapped in nearly-harmonic potentials, such as optical tweezers, is quantized [217, 218], and acts essentially as a bosonic degree of freedom, similarly to auxiliary bosonic states in ion traps [219, 220] and superconducting circuits [221, 222]. Recent experiments exploiting this resource have demonstrated squeezing of motional states [217] and creation of a quantum register of fermion pairs [218]. Unlike electronic or nuclear states, motional states do not interact directly with electromagnetic fields and are thus robust to environmental effects [223]. Further, their bosonic nature may be used to implement quantum error correction schemes with bosonic degrees of freedom such as Gottesman-Kitaev-Preskill (GKP) codes [220, 224], or as a resource for studying lattice gauge theories [225, 226].

Building from the efficient preparation of arrays of motional ground state atoms in Section 3.5, we demonstrate key tasks in quantum information processing using motional states. We first demonstrate coherent transduction of electronic superposition states into the two lowest energy motional states along a single tweezer axis, with a one-way fidelity of $0.993(3)$, post-selected on erasure detection. We then combine this with local addressing via selective trap depth control [33] in order to perform mid-circuit readout of a tweezer array [106, 139, 196, 197, 223]. To our knowledge, this is the first such demonstration on an optical transition; furthermore, it relies only on local tweezer depth control, and does not require multiple species [223], nuclear states [106, 197], or a cavity [139]. Importantly, as part of this study, we also demonstrate mid-circuit erasure detection [75], utilizing metastable motional state superpositions which have a coherence time much longer than our mid-circuit detection time.

Assuming perfectly harmonic trapping, the manifold of motional states for a single principal axis acts as a bosonic degree of freedom with an infinite ladder of states. We will focus on the case of controlling only a single motional degree of freedom, and in particular focus

on the lowest two levels, which we refer to as the *motional qubit*. The isolation of these two motional states is achieved due to driving sideband transitions which only couple the motional ground state to the first excited state, and therefore higher energy levels are not populated. Extensions of our scheme and results are readily generalized to control over all three axes, and to control over higher energy states in the motional manifold, which may enable various new applications [220] not explored directly in this work. We will use a trap frequency (equivalently, motional qubit energy spacing) of $\omega/(2\pi) \simeq 35.5$ kHz.

4.2.1 Preparing motional superposition states

We first show the preparation of a coherent superposition state of the motional qubit via transduction from an optical qubit (Fig. 4.4a). Essentially, the preparation and readout of this state is a form of Ramsey interrogation, with the addition of spin-motion coupling inserted into the Ramsey dark time. We first perform ECC to isolate atoms in $|0, \uparrow\rangle$. We then form a spin-superposition state by performing a $\pi/2$ pulse on the carrier, and transduce this superposition onto the motional qubit by performing spin-motion coupling via sideband driving. This yields the state $|+, 0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |\uparrow\rangle$. We then leave this state unperturbed for a duration τ during which it will evolve coherently, and then read its state by reversing the sideband and carrier driving.

We plot the resultant signal, given as the probability P_{\downarrow} for the atom be in $|\downarrow\rangle$ after the entire sequence (Fig. 4.4b). We observe Ramsey oscillations at a frequency of 35.3(3) kHz, consistent with the programmed trap frequency, as is expected for the motional qubit.

Importantly, after the transduction, a coherent motional qubit state is sensitive to dephasing largely only from noise on the trapping laser *power* (which controls the relative qubit energy spacing). This is in contrast to optical qubits, which may dephase due to noise on the addressing laser *frequency*, often limiting the coherence time for optical qubits based on ultranarrow transitions [69]. We demonstrate this insensitivity by fixing $\tau = 100 \mu\text{s}$, and then vary the laser phase between each preparation or readout pulse (Fig. 4.4c). We observe that P_{\downarrow} is apparently insensitive to laser phase variation between the two sideband pulses (green points), while the atom is in a motional superposition state. By contrast, significant sensitivity is observed with respect to phase variation between the carrier and sideband pulses (red and blue points and lines), as during that time the atom is in a sensitive optical superposition state. In practice, however, the time between the carrier and sideband pulses is short enough to avoid significant dephasing.

Thus, our results show that motional levels can be used to store information in a subspace

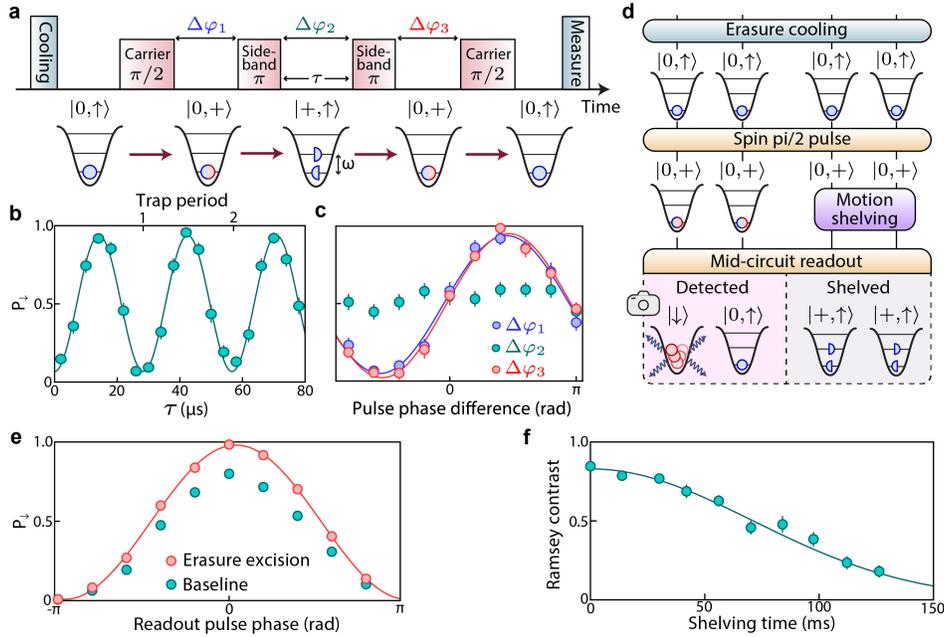


Figure 4.4: Spin-motion transduction and mid-circuit readout. **a**, Sketch of the protocol used to transfer a coherent superposition of pseudo-spin states into a superposition of motional states. After cooling the atoms and preparing the spin-superposition state $|0, +\rangle$, we apply a sideband pulse which drives the $|0, \downarrow\rangle \rightarrow |1, \uparrow\rangle$ transition, allowing us to obtain the motion-superposition state $|+, \uparrow\rangle$. We then revert the process in order to read out the coherence of the created state. **b**, Probability to obtain the atom in $|\downarrow\rangle$ as a function of holding time in $|+, \uparrow\rangle$. We observe an oscillation with frequency $35.3(3)$ kHz, consistent with our trapping frequency $\omega/(2\pi) \simeq 35.5$ kHz. **c**, Probability to obtain the atom in $|\downarrow\rangle$ as a function of the phase difference between the pulses in **a**. The atom is sensitive to the laser phase while in a pseudo-spin superposition state. However, when in the motion-superposition state $|+, \uparrow\rangle$, the atom is not sensitive to the laser phase, and is hence insensitive to laser phase noise. **d**, Sketch showing local pseudo-spin coherence preservation using motional state shelving. We prepare the $|0, +\rangle$ state and locally shelve some atoms into $|+, \uparrow\rangle$ by lowering their trap depth [33] and applying site-selective sideband drive. We detect atoms in $|\downarrow\rangle$ to project the unshelved atoms into $|\downarrow\rangle$ or $|\uparrow\rangle$. The shelved atoms are not affected by the detection, hence the coherence is preserved within the motional states. **e**, Obtained pseudo-spin coherence for the shelved atoms after the sequence described in **d**. We obtain a bare one-way transduction fidelity of $0.943(3)$, further increased to $0.993(3)$ with erasure detection and excision. **f**, Motional coherence in an echo sequence as a function of shelving time. We obtain a typical $1/e$ coherence time of ~ 100 ms, which is long compared to our fast detection time of $24 \mu\text{s}$. Error bars represent the standard error of the mean.

that is insensitive to phase noise [115] on the laser which drives the optical qubit transition. Other extrinsic noise sources, such as environmental magnetic and electric field noise should behave similarly, assuming they minimally affect the trap frequency. This overall

insensitivity makes the demonstrated motional qubit a promising storage subspace for quantum information, which we discuss in the following.

4.2.2 Mid-circuit readout with motional state shelving

Having demonstrated our capability to prepare motion-superposition states, we now show that this can be used to perform mid-circuit readout [106, 139, 196, 197, 223] (Fig. 4.4D), which is a necessary step in error correcting codes [227], and which enables various near term applications such as measurement-based generation of long-range entangled states [228, 229] and enhanced metrology [200].

To do so, we first divide the atom array into two ensembles A and B . We apply a global carrier $\pi/2$ pulse on both ensembles to prepare the $|0, +\rangle$ state. We then selectively lower the trap depths in B , while holding A constant [33]. This induces a reduction in trap frequency for atoms in B , and thus changes the motional qubit energy spacing. We can then selectively apply sideband driving only to the ensemble B , hence transferring only atoms in B into the $|+, \uparrow\rangle$ state [33], leaving A largely unperturbed.

We then apply a global, fast detection of atoms in the electronic ground state [26], which projects the spin state of atoms in A , while preserving the coherence of atoms in B . We finally unshelve atoms in B , and read out their coherence (Fig. 4.4e). The shelving and unshelving procedure is designed to apply a motional echo composed of four sideband pulses in order to eliminate shot-to-shot fluctuations of the trap depth (Fig. 4.5). We obtain a bare contrast of 0.79(1) after performing these four sideband pulses, which corresponds to a state preparation and measurement (SPAM) corrected one-way, optical-to-motion transduction fidelity of 0.943(3). This fidelity can be further improved thanks to erasure conversion of atoms in B . If the transduction fails, then atoms in B are detected during the fast image, which we then excise to obtain an overall contrast of 0.972(10) which corresponds to a one-way transduction fidelity of 0.993(3). We note that as the fast image has negligible impact on the coherence of motional states in B , this protocol demonstrates *mid-circuit* erasure detection which has only been recently demonstrated with ytterbium atoms in neutral atom systems [75].

We measure the coherence time of the motional states for atoms in B (Fig. 4.4f) using a motion-echo sequence, and obtain a typical $1/e$ time of ~ 100 ms, which is long compared to our imaging time of $24 \mu\text{s}$ using our fast detection scheme [26]. This coherence time is thought to be limited by tweezer-induced Rayleigh scattering and tweezer intensity fluctuations; we expect these results should be improvable by using better stabilization

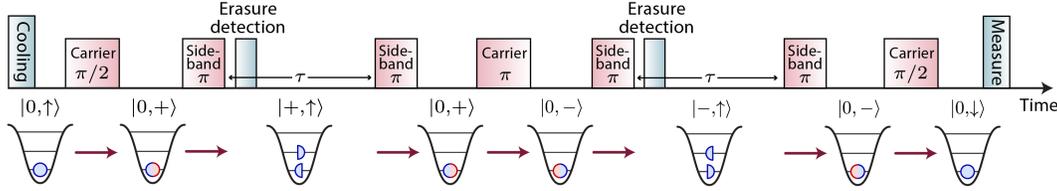


Figure 4.5: **Motional-echo sequence.** Sketch of the motional-echo sequence used in Fig. 4.4e, here assuming the phase accumulated during the wait time τ is exactly a multiple of 2π solely for presentation simplicity.

techniques, although we leave such a study to future work.

These high transduction fidelities, combined with the relatively long motional state coherence time as compared to coherent manipulation and detection duration, make this new shelving method an appealing alternative to other recently demonstrated schemes for mid-circuit detection in tweezer-based platforms [106, 139, 196, 197, 223]. As opposed to other schemes, our method does not require any specific electronic or nuclear structure, besides the presence of a sideband resolved transition, and is applicable for alkaline-earth as well as both molecular and alkali species. Further, so far the demonstrated schemes for mid-circuit readout require an auxiliary laser to address specific atoms in the array. In our case, the single-site addressability only requires single tweezer depth control [33], which is a built-in feature of most tweezer platforms. We note this is fundamentally different than methods which protect atoms from imaging light via detuned narrow-line imaging [102], and offers many distinct advantages we discuss below.

4.2.3 Motional coherence and motional-echo sequence

Here we describe the motional-echo sequence we use for the results presented in Fig. 4.4e. The sequence is similar to the one presented in Fig. 4.4a, and is further detailed in Fig. 4.5. In the absence of such an echo sequence, we find relatively short motional coherence times of up to ~ 5 ms, limited by both trap frequency inhomogeneities across the array, and shot-to-shot fluctuations in trap frequency.

The scheme proceeds as follows: (1) we prepare the motion-superposition state $|+, \uparrow\rangle$ as described in the main text. (2) The state is then allowed to evolve for a duration τ . (3) We then convert the state back into the electronic manifold, and apply a carrier π -pulse, which exchanges the populations in $|\uparrow\rangle$ and $|\downarrow\rangle$. (4) The state is then transduced back into the motional manifold, meaning that effectively we have exchanged the populations in $|0\rangle$ and $|1\rangle$. (5) The state then evolves for time τ , realizing effectively a standard spin-echo sequence. Erasure-excision is performed on this sequence by performing a fast image just

before steps (2) and (5), and post-selecting on detected atoms.

The motional-echo sequence used to obtain the results of Fig. 4.4e consists of four sideband pulses. In order to derive the one-way transduction fidelity $F_{\text{OW}} = 0.993(3)$ quoted in the main text from the bare measured contrast of $F_{\text{tot}} = 0.972(10)$, we assume that each sideband pulse is independent from the previous ones, meaning $F_{\text{tot}} = F_{\text{OW}}^4$. This assumption is reasonable as we perform erasure detection between sideband pulses, which both suppresses any potential coherent superposition between $|0, \downarrow\rangle$ and $|1, \uparrow\rangle$, and also ejects any atom left in $|\downarrow\rangle$.

Finally, we note that the fundamental limit to the motional coherence time would stem from Rayleigh scattering. This scattering rate was calculated [133] for the fermionic ^{87}Sr to be about 10^{-4} s^{-1} at a trap depth of a single recoil energy $U = E_r$; for comparison we typically operate at a trap depth of $U \approx 250E_r$ ($\approx 1300E_r$ during imaging).

4.3 ENTANGLEMENT AND HYPER-ENTANGLEMENT

As the final tool of our universal quantum processor, we demonstrate entangling gates between long-lived qubits. In particular, we will use transient Rydberg excitation to entangle two clock qubits, with applications in quantum computing and quantum metrology, as we will discuss. In this particular demonstration our entanglement fidelity was limited to roughly 90% for unimportant technical reasons, but in more recent studies we have improved the two-qubit gate fidelity to the 0.997 level [39], as measured via randomized benchmarking.

Further, we will showcase how we can combine multi-qubit entanglement with coherent, independent manipulation of electronic and motional degrees of freedom to generate a *hyper-entangled* state [202, 203, 230], i.e. a state of two atoms which is entangled across two degrees of freedom. This has the potential to enable various applications such as hyper-parallel quantum computation [202], entanglement purification [203, 231], or superdense coding [232, 233]. Further, hyper-entanglement could be used to implement two-copy interference protocols [234] using a single system, where one copy is encoded in motional states, and the other one in the electronic degree of freedom. Our scheme is naturally extensible to control over all three motional axes, allowing up to four Bell states worth of entanglement (one electronic, three motional) to be stored in a single pair of atoms. Given the long-lived nature of the motional entanglement, one could perform quantum computation tasks in the spin degree of freedom, and quantum memory tasks on the motional degree of freedom.

Building upon the single qubit control showcased in the previous sections, we now turn to generate entanglement between motional states of atoms in separate tweezers. We first demonstrate the preparation of motion-entangled Bell pairs, i.e. two atoms whose center-of-mass motion is entangled while their spin state is separable. Second, we demonstrate the simultaneous hyper-entanglement of multiple degrees of freedom per particle [202, 203]. Here, these multiple degrees of freedom are the electronic state of the atoms and their motion along a single axis.

A natural extension of our work is to use the full bosonic degree of freedom to implement quantum error correction schemes [220, 224], or to study lattice gauge theories [225, 226]. Further, the controlled entanglement of the motion of atoms separated at mesoscopic distances could lead to entanglement-enhanced tests of short-length gravitational effects [235].

4.3.1 Entanglement

We entangle the center-of-mass motion of two atoms by combining spin-motion conversion and spin-spin interactions mediated via excitation to high-lying Rydberg states (Fig. 4.6a). Starting with two atoms in $|0, \uparrow\rangle$, we first apply a spin $\pi/2$ pulse in order to prepare each in $|0, +\rangle$, for which we denote the two-atom state as $|00, ++\rangle$. We then apply a controlled-Z gate through Rydberg interactions [39, 74, 75, 95, 119]. A final $\pi/4$ pulse rotates the state into the spin-entangled Bell state $|00, \Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle \otimes (|\downarrow\downarrow\rangle + |\uparrow\uparrow\rangle))$. We then perform spin-motion transduction, obtaining a state which is spin-separable, but motion-entangled, $|\Phi^+, \uparrow\uparrow\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \otimes |\uparrow\uparrow\rangle$.

We characterize the resulting state by performing a parity measurement on the motional degree of freedom (Fig. 4.6b). This consists of evolution under the trapping potential for a variable amount of time, which induces a phase difference between different motional states, followed by spin-motion transduction for readout. We observe an oscillation frequency of $\omega/(2\pi) \simeq 71.0(5)$ kHz, consistent with the expectation of twice the trap frequency (green markers).

We also generate the odd parity Bell state [95] $|\Psi^+, \uparrow\uparrow\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) \otimes |\uparrow\uparrow\rangle$ by altering the spin- $\pi/4$ rotation to yield the spin entangled state $|00, \Psi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle \otimes (|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle))$, followed by transduction to the motional state. In this case, we do not observe a parity oscillation, since the energy gap between $|01\rangle$ and $|10\rangle$ is given only by the small difference in trap frequency between the two traps (Fig. 4.7).

Combining the contrast of the parity oscillation with the measured populations in the states $|00\rangle$ and $|11\rangle$, we obtain a motional Bell state fidelity for the state $|\Phi^+, \uparrow\uparrow\rangle$ of

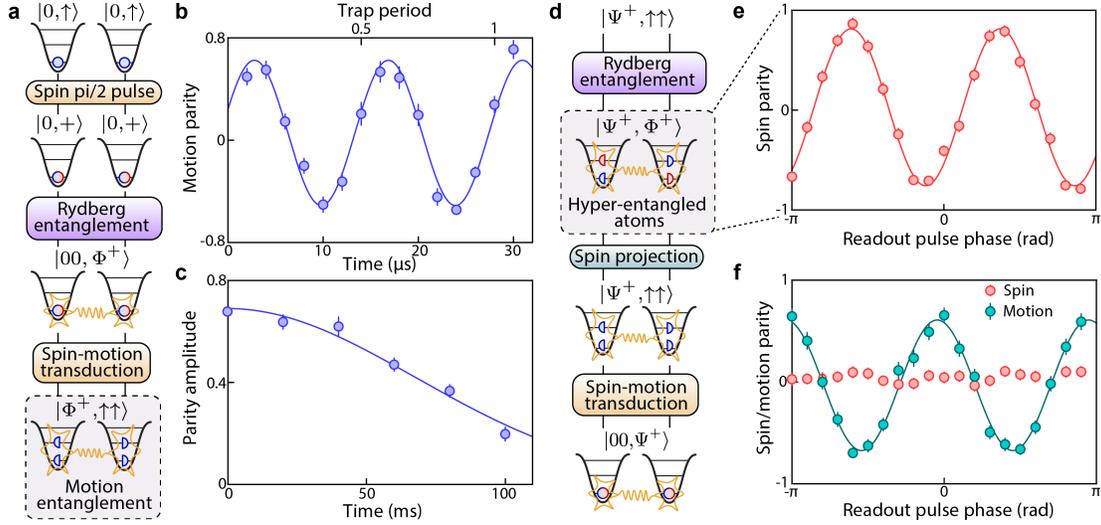


Figure 4.6: Entanglement and hyper-entanglement of motion and spin. **a**, Sketch of the protocol used to generate an entangled state of motion. After cooling the atoms and preparing the pseudo-spin superposition state $|0, +\rangle$, we entangle the spin degree of freedom via Rydberg interactions, preparing either the $|00, \Phi^+\rangle$ or the $|00, \Psi^+\rangle$ state (see text). We then apply spin-motion transduction via sideband driving, transferring spin-entanglement to motion-entanglement, yielding $|\Phi^+, \uparrow\uparrow\rangle$, or $|\Psi^+, \uparrow\uparrow\rangle$. **b**, Parity oscillation measured by varying the time in $|\Phi^+, \uparrow\uparrow\rangle$ before converting back to pseudo-spin states for readout. The obtained oscillation frequency of $71.0(5)$ kHz is consistent with twice the trap frequency $2\omega/(2\pi) \simeq 71$ kHz, as expected for $|\Phi^+, \uparrow\uparrow\rangle$. **c**, Motion parity amplitude as a function of time. We observe a Gaussian decay of fidelity with a typical $1/e$ decay time of $96(5)$ ms. **d**, Sketch of the protocol used to generate hyper-entanglement of both spin and motion. We prepare the motion-entangled state $|\Psi^+, \uparrow\uparrow\rangle$, and apply a second entangling operation on the spin degree of freedom, producing the hyper-Bell state $|\Psi^+, \Phi^+\rangle$. In order to read out this hyper-entanglement, we project the spin component onto $|\Psi^+, \uparrow\uparrow\rangle$, then convert the motional entanglement into spin entanglement. **e**, Spin parity oscillation after hyper-entangling the atoms measured by applying a readout $\pi/2$ -pulse and varying its phase. **f**, Parity oscillation after the spin projection, without (red) and with (green) spin-motion transduction. The parity oscillation collapses without spin-motion transduction, as the projection disentangles the spin part. However, the motion entanglement remains intact, which allows us to recover spin entanglement after performing the transduction.

$F_{\text{Bell}}^{\text{motion}} = 0.75(1)$ (uncorrected for SPAM errors). Currently, this value is limited by the bare Bell state generation fidelity in the electronic manifold and the combined transduction fidelity for both atoms. In more recent results [39] we have achieved an electronic Bell state fidelity of ~ 0.997 , and thus we expect the motional Bell state to have improved accordingly.

Importantly, the lifetime of motionally-entangled states is relatively long. We quantify this by measuring the amplitude of the parity oscillation signal as a function of hold time (Fig. 4.6c). We embed this measurement in a motion-echo sequence, as in the single atom

case (Fig. 4.5), to remove the effects of shot-to-shot trap variation. We observe a Gaussian decay, with a $1/e$ decay time of $96(5)$ ms. This coherence time is similar to the one obtained for the single-atom motional superposition case, and we expect this also to improve through better trap intensity stabilization.

These results demonstrate Bell pair generation of motional states in separate traps spaced by distances much larger than the spatial extent of the atomic wavefunction; in our particular implementation, traps are separated by $\approx 3.3 \mu\text{m}$, while the ground state wavefunction spread along the array axis is only ≈ 60 nm. To our knowledge, this is the first demonstration of motional Bell pair generation for neutral atoms in general, following earlier work with ions [236].

4.3.2 Hyper-entanglement

Using this versatile control we can go one step beyond pure motional entanglement, and demonstrate hyper-entanglement of motion and electronic state. In a hyper-entangled state, more than a single degree of freedom across two atoms is entangled, yielding a larger Hilbert space with various applications such as entanglement purification [203, 231], hyper-parallel quantum computation [202], or superdense coding [232, 233]. To date, such hyper-entangled states have only been realized utilizing multiple degrees of freedom in the state of a photon [203].

To accomplish this task, we prepare the motion-entangled state $|\Psi^+, \uparrow\uparrow\rangle$, and then apply a second entangling operation on the electronic degree of freedom. As opposed to hyperfine or nuclear qubits, the Rydberg drive couples nearly identically to both levels of the motional qubit, up to a detuning $\delta \simeq 2\pi \times 35$ kHz that is much smaller than the Rabi frequency $\Omega \simeq 2\pi \times 3.5$ MHz applied. The gate thus acts on the electronic degree of freedom, leaving the motional state largely unchanged, generating the hyper-Bell state $|\Psi^+, \Phi^+\rangle$.

To show that entanglement coexists on the two degrees of freedom, we first measure the spin parity and population (Fig. 4.6e). We obtain an entanglement fidelity $F_{\text{Bell}}^{\text{spin}} = 0.855(7)$. To show the simultaneous motional entanglement, we perform our fast imaging procedure [26] to project the $|\Psi^+, \Phi^+\rangle$ state onto $|\Psi^+, \uparrow\uparrow\rangle$ or $|\Psi^+, \downarrow\downarrow\rangle$ with equal probability. Atoms in $|\downarrow\rangle$ have a low probability of surviving the fast imaging [26], so in the following we only consider pairs projected into $|\Psi^+, \uparrow\uparrow\rangle$.

We first verify that no spin entanglement remains by measuring a null spin parity (Fig. 4.6f, red markers). We then apply spin-motion transduction, and again recover a parity oscillation signal (green markers), indicating that the motion-entanglement was preserved. We obtain

$F_{\text{Bell}}^{\text{motion}} = 0.76(1)$. Even with the complications of the hyper-entanglement sequence, this value is comparable to the one obtained using the sequence described in Fig. 4.6a, indicating that Rydberg entanglement and spin projection do not appreciably affect motion-entanglement.

4.3.3 Details of the two-qubit Rydberg gate

The controlled-Z (CZ) gate is realized by a global excitation in the Rydberg blockade regime with a time-optimal pulse [74, 75, 237, 238]. We use a single-photon excitation to a Rydberg state ($5s61s^3S_1$) of two atoms spaced by $3.25 \mu\text{m}$ with a phase modulated field. A 500G magnetic field perpendicular to the $|\downarrow\rangle \rightarrow |\uparrow\rangle$ laser beam propagation axis is applied throughout the sequence, enabling high Rabi frequency on this transition. This high magnetic field results in a reduction of the Rydberg interaction strength as compared to low field conditions [29], and we leave an investigation of this effect to future work. We measure a C_6 coefficient of $52 \text{ GHz}\cdot\mu\text{m}^6$ and use a maximal Rabi frequency of $2\pi \times 3.6 \text{ MHz}$.

For $|\Phi^+\rangle$, we measure SPAM uncorrected populations $P_{\uparrow\uparrow} + P_{\downarrow\downarrow} = 0.97(2)$; and induce parity oscillations with a final $\pi/2$ analyzing pulse where we measure a contrast of $0.89(1)$, yielding a SPAM-uncorrected optical Bell state [119] generation fidelity of $0.93(2)$. We note that following further optimization [39], we observe a SPAM-corrected controlled-Z fidelity $0.9935(10)$ in a randomized benchmarking [74] sequence, which we have since improved to the 0.997 level, though do not show those results explicitly here.

4.3.4 Parity oscillations for the motional Bell states

Here we show our results concerning parity oscillations of the odd-parity and even-parity motion-Bell states $|\Psi^+, \uparrow\uparrow\rangle = (|01\rangle + |10\rangle)/\sqrt{2} \otimes |\uparrow\uparrow\rangle$ and $|\Phi^+, \uparrow\uparrow\rangle = (|00\rangle + |11\rangle)/\sqrt{2} \otimes |\uparrow\uparrow\rangle$. After preparing these states, we perform the experimental sequence described in Fig. 4.7a, which consists of holding atoms in the motion Bell state for a variable amount of time, and then reading out the parity signal. We show the results of this experiment in Fig. 4.7b, for both $|\Phi^+, \uparrow\uparrow\rangle$ (blue circles) and $|\Psi^+, \uparrow\uparrow\rangle$ (green circles). We observe a significant difference in their behavior: for $|\Phi^+, \uparrow\uparrow\rangle$, we obtain a parity oscillation (also shown in Fig. 4.6b), whereas for $|\Psi^+, \uparrow\uparrow\rangle$, we do not observe significant parity oscillations.

This behavior is expected: as both tweezers have approximately the same power and same trap frequency ω , the energy spacing $\hbar\omega$ of the motional level is the same for both tweezers.

In order to prove that we undoubtedly prepare $|\Psi^+, \uparrow\uparrow\rangle$, we induce parity oscillations by varying the power of a single tweezer in the pair. The experimental sequence is described

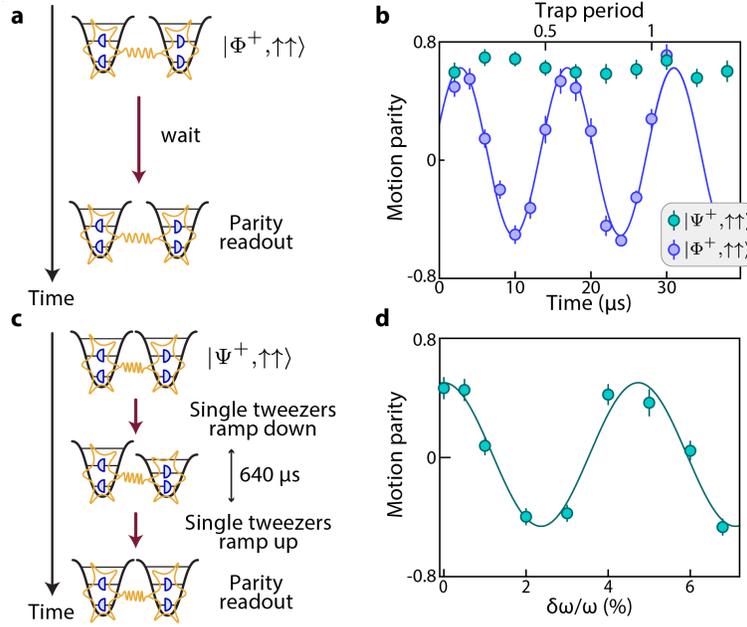


Figure 4.7: **Inducing parity oscillations on an odd-parity Bell state with relative trap depth variation** **a**, Sketch of the performed experiment to induce the parity oscillation on $|\Phi^+, \uparrow\uparrow\rangle$. **b**, Parity oscillations following the experiment described in **a** for $|\Phi^+, \uparrow\uparrow\rangle$ (also shown in Fig. 4.6b), and for $|\Psi^+, \uparrow\uparrow\rangle$, for which we do not observe significant oscillations. **c**, Sketch of the performed experiment to induce parity oscillations on $|\Psi^+, \uparrow\uparrow\rangle$. After preparing the motion-entangled Bell state $|\Psi^+, \uparrow\uparrow\rangle$, we ramp down the power of a single tweezer for a fixed duration of $640 \mu\text{s}$. **d**, Motion parity as a function of the relative trap depth difference between the two tweezers.

in Fig. 4.7c. After preparing the motion-entangled Bell state $|\Psi^+, \uparrow\uparrow\rangle$, we ramp down the power of a single tweezer for a fixed duration of $t_{\text{wait}} = 640 \mu\text{s}$. During this time, the state evolves as $(|01\rangle + e^{i\delta\omega t_{\text{wait}}}|10\rangle)/\sqrt{2} \otimes |\uparrow\uparrow\rangle$, where $\delta\omega$ is the trap frequency difference between the two tweezers.

After this step, we ramp up the tweezer power to its original value, and read-out the state's parity. We repeat this for various trap depth, and show our results as a function of $\delta\omega/\omega$ (Fig. 4.7b). We obtain a parity oscillation, indicating that we indeed prepared the $|\Psi^+, \uparrow\uparrow\rangle$ state.

4.4 APPLICATION: METROLOGY

We now demonstrate that access to robust, high-fidelity, single-site operations (Section 4.1) can enable enhanced sensing protocols for entanglement-free metrology. Further, we discuss how these protocols can be further improved through the use of mid-circuit measurement

(Section 4.2) or entanglement (Section 4.3).

Sensors based on quantum probes provide some of the most precise measurements in science [198, 239–242]. For many such systems, fundamental sensitivity limits can be improved through entanglement [164, 243–245], but in the presence of noise, a practical advantage of such schemes is not guaranteed [246, 247]. A complementary approach studies optimal metrology with entanglement-free quantum control and readout methods. For both approaches, an important figure of merit is not just the sensitivity to a given observable, but also the dynamic range over which that observable can unambiguously be estimated [200, 201, 248].

In the particular case of optical atomic clocks [198], the observable of interest is the stochastically evolving phase of a laser, θ , acting as a local oscillator, which is mapped into population imbalance of an ultra-narrow optical transition. The clock stability improves with the interrogation time, but the phase can only be unambiguously mapped when it is in the range $[-\pi/2, \pi/2]$; phases outside of this range lead to phase-slip errors, which limits the attainable interrogation time at a given phase-slip error probability in the case of local oscillator limited clocks. Optimal readout schemes [200, 201, 249, 250] could improve the attainable interrogation time exponentially but require local rotational control over sub-ensembles during the sensing protocol or local mid-circuit readout and reset, both of which have not been demonstrated prior to our work.

4.4.1 Quadrature readout

Here we show one such proposal [250] experimentally, by splitting the array into two sub-ensembles using local addressing to perform Ramsey interferometry simultaneously in two orthogonal bases, X and Y , yielding populations $P^{(x)}$ and $P^{(y)}$. While readout in a single basis limits the invertible phase range to $\theta \in [-\pi/2, \pi/2]$, readout in both bases allows this range to be extended unambiguously to $[-\pi, \pi]$ (Fig. 4.8a). Consequently, we can afford a longer Ramsey interrogation time before θ drifts outside of the invertible range, which would cause a phase-slip error. Note that while the atom number in each quadrature has been halved, this typically does not increase [200, 250] the quantum projection noise (QPN) [251] from the dual-quadrature measurement compared to a single-basis measurement (Fig. 4.12).

To implement this dual-quadrature readout, we perform Ramsey interferometry with the addition of a $\hat{Z}(\pi/2)$ rotation to all odd sites in the array before readout (Fig. 4.8b). The resultant oscillations in $P^{(x)}$ and $P^{(y)}$ show a $\pi/2$ phase shift between the even (X) and odd (Y) sites in the array (Fig. 4.8c). For every repeated measurement (indexed by j) at time t

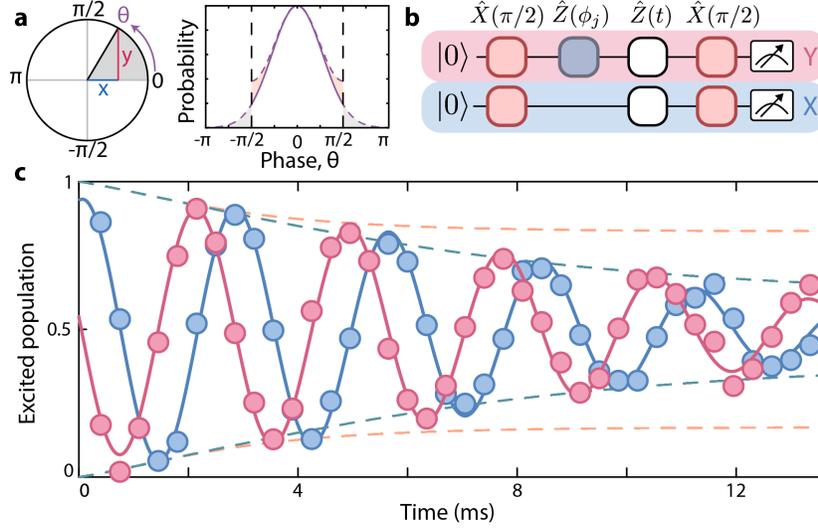


Figure 4.8: **Enhanced sensing with dual-quadrature measurement.** **a.** For a given phase angle, θ , population measurement in only a single basis, e.g. Y , can only be inverted within a dynamic range of $-\pi/2 < \theta < \pi/2$. By measuring both quadratures, X and Y , this dynamic range can be doubled to $-\pi < \theta < \pi$, allowing for interrogating larger spreads in phase, such as when measuring for longer times. **b.** We implement dual-quadrature readout of Ramsey interrogation by applying local $\pi/2$ phase shifts to all odd sites in the array. **c.** With single-quadrature readout, the interrogation time is limited due to phase-slips, visible by the separation between a decay envelope reconstructed from the single-quadrature phase spread (orange dashed line) and the averaged Ramsey signal (blue and red markers and lines). The equivalent reconstruction with dual-quadrature readout (green dashed line) is accurate up to longer times, as will be explained in Fig. 4.9.

we estimate the phase as [200]

$$\theta_j(t) = \arg(z_j^{(x)}(t) + iz_j^{(y)}(t)), \quad (4.2)$$

where $z_j^{(x,y)}(t) = (2P_j^{(x,y)}(t) - 1)$ and \arg is the argument function. We then calculate the difference, $\delta_j(t)$, of $\theta_j(t)$ from its mean phase (Sec. 4.4.4).

We plot the probability distribution $\mathcal{P}(\delta_j(t))$ in Fig. 4.9a, and observe a continuous growth of its standard deviation (STD) σ (inset). We stress that we are interested in the distribution of the laser phase itself, which determines the phase-slip error probability. Hence, we have subtracted off the contribution from QPN to our experimental data shown in the inset of Fig. 4.9a (Sec. 4.4.3). We find that the laser phase STD grows with time as a power law $\sigma = \beta t^\alpha$, with $\alpha = 0.56(2)$ which we attribute to a power spectral density composed of $1/f$ and white frequency noise. If this standard deviation of the laser phase itself becomes too large compared to the dynamic range, frequent phase-slip errors occur. In Fig. 4.9b we evaluate the phase-slip probability, ϵ , that the phase has exceeded the bounds $[-\pi/2, \pi/2]$ (in emulation of a theoretical single-basis measurement, black dashed lines in Fig. 4.9a),

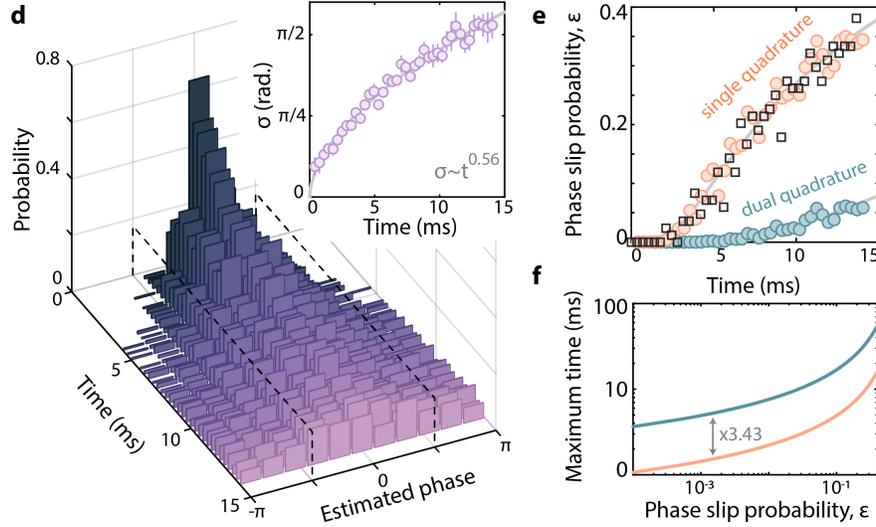


Figure 4.9: **Gains from multi-ensemble metrology.** **a.** To perform phase reconstruction, we measure time-resolved probability distributions of the estimated phase relative to the mean from dual-quadrature measurement. As the standard deviation, σ , of the phase distribution grows (inset), the estimated phase begins exceeding the $-\pi/2 < \theta < \pi/2$ range for normal spectroscopy (black dashed lines), but is still resolvable via dual-quadrature measurement. Note that the time-dependent contribution from quantum projection noise to the standard deviation has been subtracted off in the inset (Sec. 4.4.3). **b.** We estimate the phase-slip probability, ϵ , for single- (orange circles) and dual-quadrature (green circles) measurements by fitting a folded Gaussian to the time-resolved estimated phases in **a**. The fit is folded over at the boundaries of the dynamic range to account for the behavior of phase-slips, as in Fig. 4.8a. For the single-quadrature case, we also estimate the probability directly from the underlying data (squares), which is in good agreement with the estimate from fit. Solid lines are the predicted phase slip probabilities from the fit in the inset of **a**. This fit is used to estimate decay envelopes in Fig. 4.8c. **c.** For a given allowable phase-slip probability, the enhanced dynamic range of the dual-quadrature readout improves the maximum possible interrogation time. For our particular phase growth profile (inset of **a**), the improvement is a factor of ~ 3.43 .

or $[-\pi, \pi]$ (for the dual-quadrature readout); we find that the error probability for the single-basis case quickly becomes substantially larger at shorter interrogation times.

We characterize the maximum interrogation time, $T_{\max}(\epsilon)$, for which the phase-slip error probability is still below a threshold ϵ (Sec 4.4.4). We find that $T_{\max}(\epsilon)$ is significantly increased for the dual-quadrature case (Fig. 4.9c) by a factor of 3.43(13), as determined by the phase STD growth rate observed experimentally and related to the laser noise spectrum. Such elongation in the attainable interrogation time can be translated directly to enhanced stability in a metrological setting. For example, in a zero dead-time optical clock the stability is proportional to the square-root of $T_{\max}(\epsilon)$, such that we can project an increase in stability by a factor of $\sqrt{3.43} \sim 1.8$ for our particular noise profile. This would constitute a

practical improvement in phase estimation without increasing the probability of phase-slip errors, a common problem for entanglement enhanced metrology schemes [199, 247].

4.4.2 Multi-ensemble metrology

Even greater enhancements in dynamic range, and hence clock stability, could be possible through the use of multiple ensembles with different interrogation times by utilizing fast QND measurements [252, 253] or by explicitly programming ensembles with different sensitivities to the global laser phase [200, 201]. In the latter of these protocols, the total number of atoms is evenly divided into M ensembles, which are each further subdivided into two sub-ensembles for dual-quadrature measurement. One ensemble is used for normal phase measurement, while for the rest the free evolution time is reduced by factors of $2^{-1}, \dots, 2^{1-M}$, or equivalently their effective phase accumulation is reduced by the same amount. If this procedure is performed correctly, the effective ensemble coherence times will then be extended by factors of $2, \dots, 2^{M-1}$, meaning slower evolving ensembles can be used to probe for phase-slips in the fastest ensembles. This then allows for phase estimation over a wider dynamic range beyond $[-\pi, \pi]$, and potentially allows for an improved scaling of the clock stability with atom number [200] at fixed phase-slip probability (Fig. 4.10a). Such a scheme can be realized through several physical mechanisms, including local mid-circuit readout, local dynamical decoupling, or by preparing a cascade of progressively larger GHZ states [199].

To give a simplified analogy, the idea is basically the same as having *multiple hands on a watch*: an hour hand to keep track of longer time changes, and a minute hand to more precisely keep track of shorter time changes. What we have done is essentially build such a multi-hand clock at the atomic scale. In this analogy, by adding more ensembles, we add more hands to our watch which keep track of longer (days, weeks, months, years, etc.) or shorter (seconds, milliseconds, etc.) timescales.

As an outlook, we demonstrate a proof-of-principle of local control techniques towards such protocols by performing local dynamical decoupling such that three ensembles experience different effective Ramsey evolution times of $T, T/2$, and $T/4$. This is accomplished by inserting local $\hat{X}(\pi)$ pulses (using techniques from Fig. 4.1) during the evolution at time $T/4$ for the second-fastest ensemble, and time $3T/8$ for the slowest ensemble. Each ensemble is then subdivided further into two sub-ensembles for dual-quadrature readout (Fig. 4.10c). Resultant Ramsey oscillations versus the total evolution time, T , show a frequency ratio of 1:1.99(1):4.10(4), very close to the desired 1:2:4 ratio.

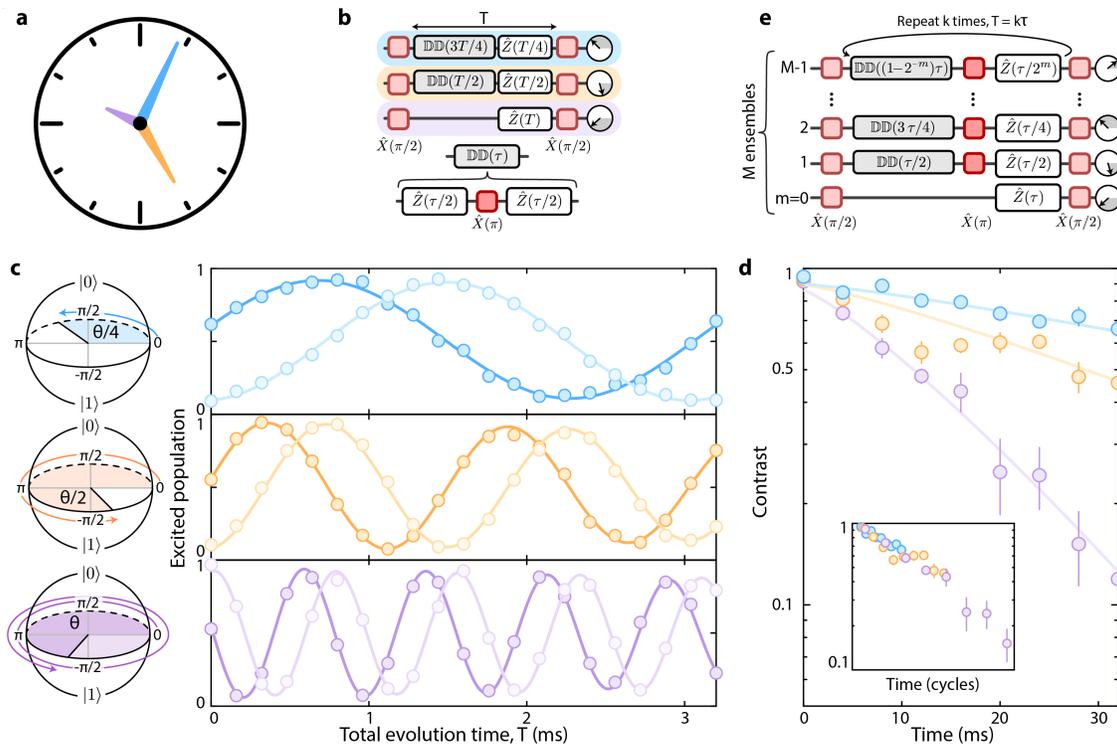


Figure 4.10: Local dynamical decoupling towards optimal metrology. **a.** Clock performance can be improved on all time-scales by using multiple *hands*. **b.** We split the array into three ensembles, and perform a local dynamical decoupling (DD) sequence such that even though the total Ramsey dark time is T , individual ensembles experience different effective evolution times of $T/4$, $T/2$, and T , respectively. The phase of each ensemble is then measured using dual-quadrature readout. **c.** Slower evolving ensembles (those which experience less evolution time) can be used to detect phase-slips in faster evolving ensembles, extending the effective interrogation time of optical clocks. Following the sequence in **b**, we find the three ensembles evolve at relative rates of 1:1.99(1):4.10(4) with respect to the total evolution time, T . **d.** For a slightly different scheme (not shown), we measure the contrast of the different ensembles over longer evolution times, which we find scales with the phase accumulation rate (i.e. the decoherence goes as the number of cycles, as expected). The demonstrated scheme in **b-d** is effective for the case of slow frequency noise where the corresponding noise correlation time is longer than the total evolution time. **e.** To handle generic time-dependent noise with shorter correlation times, we envision breaking the total evolution time into k kernels of length τ , each of which is composed of local dynamical decoupling and free evolution. In this way, as long as τ is shorter than the correlation time of any time-dependent noise affecting the system, the different M ensembles (indexed by $m = 0, \dots, M - 1$) can accumulate phase in a correlated manner over the interleaved Ramsey interrogation periods.

Following this experimental demonstration, we now discuss two limitations (and possible solutions) of this scheme, specifically related to the frequency noise profile and the atom number per ensemble. First, for the simplest case of shot-to-shot noise of laser frequencies

that are otherwise constant during the interrogation, our scheme would allow the clock stability to be improved exponentially [200] by a factor of $\sqrt{2^{M-1}/M}$; the factor of $\sqrt{1/M}$ stems from increased QPN in the ensemble used for phase estimation and assumes the total number of atoms is distributed uniformly across the M ensembles. However, for more general time-dependent frequency noise, the situation is more complex, requiring a higher-order pulse sequence [200]. We propose one such pulse sequence in Fig. 4.10e, by breaking the total evolution time, T , into multiple kernels of length τ . Within each kernel, each ensemble experiences a combination of local dynamical decoupling and free evolution, such that the net phase evolution time is $T, T/2, \dots, T/2^{M-1}$. This scheme could handle noise profiles where the local phase accumulation period, τ , is shorter than the correlation time of the noise. We numerically find that exponential scaling of the maximal interrogation time is then possible up to a saturation point set by the effective decoupling bandwidth (Fig. 4.13).

Second, multi-ensemble estimation schemes in general require sufficient atom number per ensemble to be useful [200]. When the number of atoms per ensemble is limited, quantum projection noise can negate any advantage, i.e. when the error probability in estimating a phase slip by using a slower evolving ensemble exceeds the actual phase slip probability in the fastest evolving ensemble. For the present experimental demonstration with $N \approx 6$ per ensemble we do not expect a metrological gain (Fig. 4.13), but we note that a generalization of our addressing scheme to two dimensional tweezer clock systems [186] is straightforward. For example, we imagine a realistic scenario of a two-dimensional atom array with column-by-column control of tweezer positions, such as could be generated with crossed AODs or an AOD combined with a spatial light modulator. In this case, each pair of columns could realize one ensemble with dual-quadrature readout. Finally, we note that such exponential scaling is possible only up to a time-scale where decoherence is dominated by local oscillator noise. Beyond that, interrogation time will be limited by atomic coherence and ultimately by atomic state lifetime [110].

In summary, we have used local rotations to interrogate two atomic ensembles simultaneously for dual-quadrature readout of a Ramsey interferometry signal with demonstrable metrological gain. As an aside, we note that our local phase control may also find use for mapping spatial phase variations (such as Gouy phase) across the clock interrogation beam, as can be limiting in some metrological applications like compact clocks [254]. More importantly, we have extended our dual-quadrature work to show a proof-of-principle for controlling many ensembles with variable sensitivity during Ramsey evolution, a key ingredient of proposals for optimal clocks. Further, these methods could be naturally combined

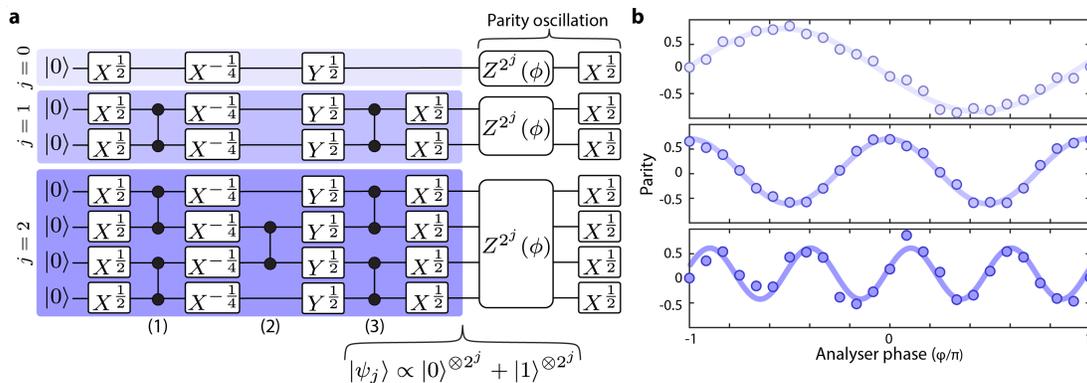


Figure 4.11: **Simultaneous interrogation with a cascade of GHZ states a.** A quantum circuit for the simultaneous preparation and characterization of GHZ states containing 2^j atoms each, here $j = 0, 1, 2$. **c.** Parity oscillations for $j = 0, 1, 2$ from top to bottom showing a 1:1.96(6):3.9(1) ratio in phase sensitivity for states of the form $|\psi_j\rangle \propto |0\rangle^{\otimes 2^j} + |1\rangle^{\otimes 2^j}$.

with metrologically useful entangled states [199, 250, 255] to enable simultaneously high sensitivity with a large dynamic range. More generally, our results are an important step towards a fully programmable quantum optical clock based on neutral atoms, which would incorporate quantum computing techniques towards metrological gains, similar to work done with ion trap devices [255] but likely in a more scalable fashion. Such a universal neutral atom clock system would ideally combine arbitrary local rotations, as shown here, with two-qubit entangling operations for optical transitions [39, 119].

As a first step towards this goal, we utilize the two-qubit entangling gates discussed in Section 4.3 to realize a cascade of progressively sized GHZ states, which can be used [199] as the different ensembles in the multi-ensemble metrology schemes [200, 201]. These results are shown in Fig. 4.11, and expanded upon in great detail in Ref. [39]. In particular we simultaneously realize three different GHZ states, with 1, 2, and 4 atoms respectively, utilizing a combination of local and global rotations, along with dynamic reconfigurability [37] to change the connectivity of two-qubit gates. The fidelity of preparing these states is limited by the coherence of our clock laser, and could be substantially improved [39, 136]. Still, these results demonstrate that our multi-ensemble schemes can be combined with entanglement; further improvements can be made by performing quantum logic spectroscopy with entangled-ancilla atoms in a mid-circuit measurement in order to efficiently prepare, evolve, and readout the sensing capabilities of atomic clocks [39].

4.4.3 Estimating laser phase prior width

Here, we detail how we find and isolate the prior width of the laser phase distribution from the experimentally measured values. We note that while the phase slip probability depends only on this prior width, the measured distribution is further affected by quantum projection noise (QPN). The QPN itself is a function of the fraction of excited atoms measured in the Ramsey sequence. One thus finds that the relative contribution of the QPN term varies with the central phase of the laser and the prior width of the phase distribution. Assuming a system of $N = 20$ atoms, we plot the calculated phase distribution width as a function of central laser phase $\bar{\theta}$, but with the underlying prior phase distribution having a width of zero (Fig. 4.12a). We repeat this calculation for both single and dual quadrature phase estimation.

We note that these estimators are affected slightly differently by QPN. Specifically, while the optimal single quadrature working point in terms of minimal QPN is around a mean phase of 0 (which corresponds to measurement with an excitation fraction of 0.5), we find that the optimal working point for a dual-quadrature estimation is around a mean phase of $\pi/4$ (corresponding to the two quadratures having excitation fractions of 0.85), though we note the dual-quadrature value is relatively flat over the entire bandwidth. For longer interrogation times, the prior width grows as a power law which depends on the laser noise spectrum. The contribution from projection noise is thus in general time-dependent.

To isolate this effect and learn the true laser phase distribution as a function of interrogation time, we first calculate the total observable width including QPN, σ_{tot} , for a range of prior widths σ_δ at a given laser central phase $\bar{\theta}$. This is done by sampling random phases from a normal distribution, followed by sampling the observed phase from a binomial random process representing the projection uncertainty. Repeating this process over a million draws we obtain the observed distribution as a function of the prior width (Fig. 4.12b). We then invert the function to obtain $\sigma_\delta(\sigma_{tot})$ and interpolate the latter to find the prior width at the given measured σ_{tot} . While the details of added noise are important for the attainable stability, we note that quantum projection noise itself can not cause a phase slip and thus does not limit the interrogation time.

4.4.4 Fitting the phase-slip probability

To find the phase deviation from the mean phase for a given shot, we fit the Ramsey oscillations in Fig. 4.8c with a decaying sinusoid, and at each time define the fitted mean populations as $\bar{P}^{(x)}(t)$ and $\bar{P}^{(y)}(t)$. These populations are inverted via Eq. 4.2 into the mean phase, $\bar{\theta}(t)$,

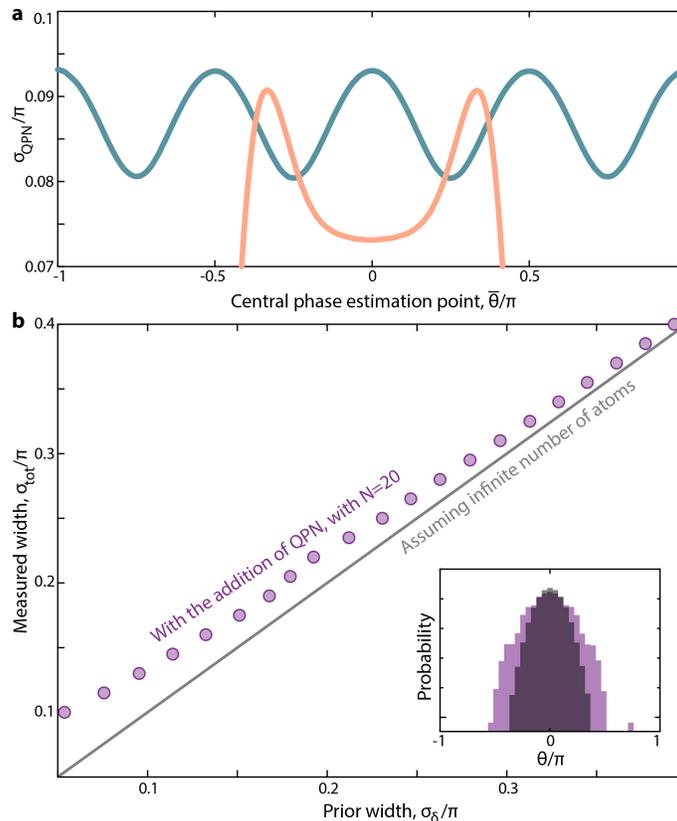


Figure 4.12: **Quantum projection noise in a dual-quadrature measurement.** **a**, Added standard deviation due to quantum projection noise (QPN) for phase estimation around different average phases $\bar{\theta}$, plotted for $N = 20$ atoms utilized in a single-quadrature (orange) or dual-quadrature (green, 10 atoms per quadrature) measurement. The added QPN varies with the phase the measurement is taken at; thus as the prior width of the phase distribution grows over time, and a broader range of phases is sampled, the QPN will vary. **b**, To learn the prior width from the measured width we sample random phases from a normal distribution, followed by sampling the observed phase from a binomial random process representing the projection uncertainty (inset). We use the sampled distributions for the dual-quadrature estimator to calculate the width including QPN for a range of prior laser widths. We then invert this function and interpolate if needed, to find the prior width for any measured width.

and finally we calculate the phase deviation from the mean as $\delta_j(t) = \text{mod}(\theta_j(t) - \bar{\theta}(t), \pi)$.

For the phase-slip probability (Fig. 4.9b), we fit the probability densities $\mathcal{P}(\delta_j(t))$ with $G(\mathcal{P})$, where G is a Gaussian distribution folded into the range $[-\pi, \pi]$. This fit provides an estimation of the true standard deviation, σ , of the $\delta_j(t)$ distribution. With this in hand, for a given half-dynamic range, B , we then find the phase-slip probability, ϵ , as

$$\epsilon = 2 \int_B^\infty G(\mathcal{P}) d\mathcal{P} = \text{erfc}\left(\frac{B}{\sqrt{2}\sigma}\right), \quad (4.3)$$

where erfc is the complementary error function. Note that here $B = \pi/2$ corresponds to a single-basis measurement while $B = \pi$ corresponds to a dual-quadrature measurement.

In order to calculate the maximal interrogation time, we first find the laser phase prior width (as described in the previous section) and then fit the time-resolved profile of $\sigma(t)$, as $\sigma(t) = \beta t^\alpha$. We find $\beta = \pi \times 0.119(6)$ and $\alpha = 0.56(2)$. The growth of σ over time can then be used to predict the Ramsey decay envelope, C , for different choices of B , as $C = e^{-\sigma^2/2}$. In Fig. 4.8c, we show this envelope estimation for the cases of $B = \pi/2$ and $B = \pi$ (orange and green dashed lines, respectively).

We then analytically calculate the maximum interrogation time at a fixed phase-slip error probability, ϵ , from Eq. 4.3 as

$$T_{\max}(\epsilon) = (B/(\sqrt{2}\beta\text{erfc}^{-1}(\epsilon)))^{1/\alpha}. \quad (4.4)$$

4.4.5 Limits in multi-ensemble metrology

To study the possible limitations of the multi-ensemble scheme we simulate stochastic phase evolution of a local oscillator with $1/f$ frequency noise, whose overall power sets a characteristic single ensemble $1/e$ Ramsey coherence time T_{LO} . We numerically find the maximal interrogation time at a fixed phase slip probability ϵ (here we use $\epsilon = 5 \cdot 10^{-3}$) with increasing ensemble number M by iteratively correcting for phase slips as described in Ref. [200]. We repeat this calculation for different dynamical decoupling block lengths τ . In Fig. 4.13a we plot the results for up to $M = 9$ ensembles, assuming infinite atom number. We find that adding more ensembles indeed enables exponential scaling of the interrogation time up to a saturation point set by the effective dynamical decoupling bandwidth (expressed in terms of τ/T_{LO}).

We further study the effect of quantum projection noise on the efficacy of the scheme in the case of low atom number per ensemble. For the optimal dynamical decoupling sequence found previously, we vary the number of atoms per ensemble N and repeat the calculation, which is now affected by quantum projection noise. Specifically, the use of slower evolving ensembles with limited atom number for the iterative correction of phase slips in the fastest evolving ensembles is prone to errors due to the increased variance in such estimation. For a small number of atoms per ensemble, this negates any advantage. However, we note that $N \simeq 20$ atoms per ensemble suffice for efficient operation of the scheme, with the interrogation times and noise strength tested here.

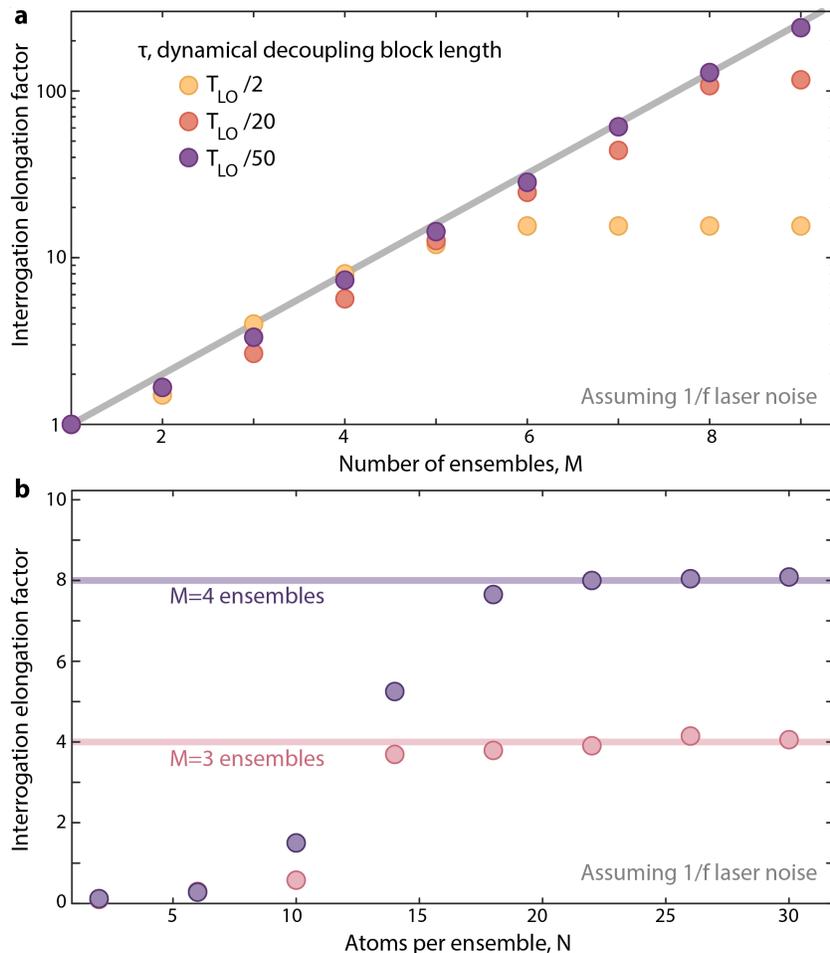


Figure 4.13: Limits in multi-ensemble metrology. a, Asymptotic scaling of the extended interrogation time factor as a function of the number of ensembles M employed. We numerically calculate the maximal interrogation time at a fixed phase slip probability for different dynamical decoupling block lengths τ , in terms of the local-oscillator coherence time T_{LO} , assuming $1/f$ frequency noise and infinite atom number. The addition of more ensembles enables exponential scaling of the maximal interrogation time (solid line marks 2^{M-1}) up to a saturation point set by the effective decoupling bandwidth. The latter can be extended by reducing the block length while maintaining a sufficiently high Rabi frequency with respect to the fast noise frequency. **b**, For the optimal decoupling sequence found in **a**, we plot the extended interrogation time as a function of the number of atoms per ensemble N . For a small number of atoms per ensemble, quantum projection noise results in an enhanced rate of false positive indication of a phase slip, negating any advantage. We find that $N \simeq 20$ atoms per ensemble suffice for efficient operation.

Coda

Rydberg atom arrays are relatively new but are already impressively capable. Importantly, most of their error sources are currently technical, and not nearing fundamental limitations. Indeed, even apparently intrinsic error sources such as spontaneous decay can be alleviated either via advanced experimental design [86] or via advanced error correction mechanisms, as we have demonstrated [26]. As systems mature further in coming years, it is likely higher fidelity operations will be achieved throughout the entire feature stack.

Our effort to understand the error sources affecting Rydberg atom arrays has allowed us to build a robust *ab initio* error model for our system. Predictions from this error model have been accurate to the experiment [26, 29], and have allowed us to identify the most meaningful error sources in our system. While here we have tested this error model (mostly) on the level of a few qubits, for the rest of this thesis we will continue applying it in a roughly unchanged form to study and understand the behavior of many-body quantum operations. Such careful calibration of an error model has been, and will continue to be, instrumental in improving quantum systems more broadly going forward, and in many ways the major task of the NISQ era is to build better models and intuition for decoherence.

Further, we showed how these errors can be mitigated through the first experimental demonstration of erasure conversion, with significant implications for quantum technology. For Rydberg atom arrays in particular, our current demonstration of ~ 0.999 SPAM-corrected two-qubit fidelity shows Rydberg atom arrays can be equipped with entangling operations at competitive fidelities to other quantum platforms, and our modeling implies that values of ~ 0.9997 could be possible with laser noise improvements alone. Further, utilizing a cryogenic environment could freeze out blackbody decay to a large degree [86], with remaining decay detected as an erasure, leaving almost no intrinsic decoherence. Our techniques could be combined with a long-lived qubit which is dark to the fast image (as since realized in Refs. [32, 39, 75], e.g. using the 3P_0 nuclear qubit in neutral Sr [256]

and Yb [104, 226], or $S_{1/2}$ in Ca^+ and Ba^+ ions [215]. Similarly, schemes for implementing erasure conversion in superconducting circuits have been put forward [257, 258] and demonstrated [259, 260]. Such techniques could lead to drastically reduced quantum error-correction thresholds [68, 261] for fault-tolerant quantum computing.

For our part, we have further demonstrated a universal toolbox for quantum computation with Rydberg atom arrays, including arbitrary single-site control [28], two-qubit gates [26] (with fidelities since improved to the ≈ 0.997 level [39]), mid-circuit measurement [26], and built-in error mitigation through erasure correction [26, 32]. We since have also incorporated dynamical circuit reconfigurability [39] via atom shuttling [37]. Together these tools enable new applications for atom arrays in quantum computation, and are specifically useful in our tweezer clock system for near-term metrological improvements [28, 39].

More generally, the techniques we have demonstrated benefit quantum simulation [26], and could enable new paradigms for tweezer arrays that control bosonic motional degrees of freedom to serve as a quantum memory or to perform quantum information processing tasks like error correction [224]. While such schemes have been explored by ion trap quantum platforms [220], we note there are distinct advantages of the tweezer array platform. For instance, tweezers can be diabatically shifted and compressed to implement native squeezing and displacement operators on the bosonic motional manifold, and these operations can be done in parallel and arbitrarily to an array of atoms simultaneously with no cross-coupling [28]. This could enable a new paradigm for efficient error correction protocols using arrays of single physical atoms to represent arrays of logical qubits.

Part II

Randomness and verification

Prelude

Quantum mechanics is often popularly associated with randomness, mostly out of association with Born's rule. However, besides the measurement-problem, quantum science is far from random, and in fact its dynamics are unitary and deterministic. Still, many quantum algorithms and protocols developed over the last several decades presume the ability for the experimentalist to prepare random quantum states at-will, where *random* admittedly requires some definition. Applications include tasks in quantum algorithms [262–264], cryptography [265–268], and communication [269]; many further applications can be found in the first few paragraphs of Refs. [270, 271] (which are themselves good discussions of random quantum states).

However, by far one of the most major applications of preparing random quantum states is that of device benchmarking [58, 272–277], where random quantum states are used to amplify the effects of small errors into macroscopic observables. Most notably, random states were instrumental to the first declarations of quantum advantage [162]. Further, ensembles of randomly distributed, but pure quantum states are of great theoretical interest, as they for instance underlie our understanding of complexity in quantum circuits [278] and black holes [279, 280]. Thus, producing quantum states at random has become increasingly important in modern quantum science, but creating random ensembles has necessitated a high degree of spatiotemporal control [267, 274, 281–284], placing such studies out of reach for a wide class of quantum systems.

In the following chapters, we describe how we solve this problem, by predicting and experimentally observing the emergence of random state ensembles naturally under time-independent Hamiltonian dynamics. First, we provide a brief introduction to the idea of random quantum states, and then showcase how we observe such randomness in two forms: 1) from projective measurement of a larger quantum system [29, 70] (Chapter 5), or 2) from time-evolution of a globally ergodic system [27], dubbed *Hilbert-space ergodicity*

(Chapter 7). For this second case, we also make predictions — and consistent observations — for the behavior of local quantum states under time-independent evolution, the natural setting of quantum thermalization. For a closed system, where the environment is a complementary subsystem, we predict and observe a smooth quantum-to-classical transition in that observables progress from large, quantum fluctuations to small, Gaussian fluctuations as the bath size grows. This transition is universal on a quantitative level amongst a wide range of systems, including those at finite temperature, those with itinerant particles, and random circuits. Further, in a theory work [71] (not summarized in this thesis) we describe how these two perspectives are intimately linked, and point to a universal description of entropy in generic quantum systems.

However, our developments are not only related to the fundamentals of randomness; we further develop and exploit these protocols that produce random states to demonstrate benchmarking of our system through only time-independent Hamiltonian dynamics [29] (Chapter 6). We demonstrate accurate estimation of the many-body fidelity, as well as numerous applications including Hamiltonian parameter estimation, noise learning, and more. For instance, we specialize some of our theoretical results on ergodicity to the case of an open system interacting noisily with an external environment [27]. We predict the statistics of observables under largely arbitrary noise channels including those with correlated errors, allowing us to discriminate candidate error models both for continuous Hamiltonian time evolution and for digital random circuits, where we specifically use the many-body fidelity as a measurable input parameter.

Ultimately our results clarify the role of randomness and ergodicity in quantum dynamics, with fundamental and practical consequences.

Uncovering emergent randomness

Preparing random, but pure, quantum states has a number of applications [58, 262–277], but so far has been limited to systems with explicit spatiotemporal control. This is because to engineer *quantum randomness*, it has been necessary to “inject” *classical randomness*, e.g. by randomizing the circuit layout and operations. However, this puts many important applications, like device benchmarking, out of reach of a broad swathe of quantum experiments. Most importantly, it also makes studying such random behavior seemingly impossible with the most natural form of quantum dynamics, i.e. time-independent Hamiltonian evolution. Here, we refute this existing paradigm, showing that randomness *does* emerge from natural unitary quantum evolution without any explicit classical stochasticity. To do so, first we (very briefly) discuss the background of random quantum states, namely through the lens of discussing the concept of *Haar randomness*. We then show how this random behavior emerges from projective measurements of subsystems of a larger quantum state, a phenomenon intimately related to the *projected ensemble*, a term we introduce [29, 70]. In upcoming chapters we will both show how this protocol is practically useful for performing device benchmarking, as well as show how a different perspective on the same dynamics allows us to understand a quantum notion of ergodicity underlying the transition from quantum to statistical mechanics.

5.1 HAAR-RANDOM STATE ENSEMBLES

We start by illustrating the concept of pure random state ensembles via a thought experiment: consider a programmable quantum device which evolves an input state $|\Psi_0\rangle$ to an arbitrary output state $|\psi_j\rangle$ through action of unitary evolution \hat{U}_j , labeled by the program setting, j (Fig. 5.1a). If the set of states $|\psi_j\rangle$ – in the limit of many repetitions with different j

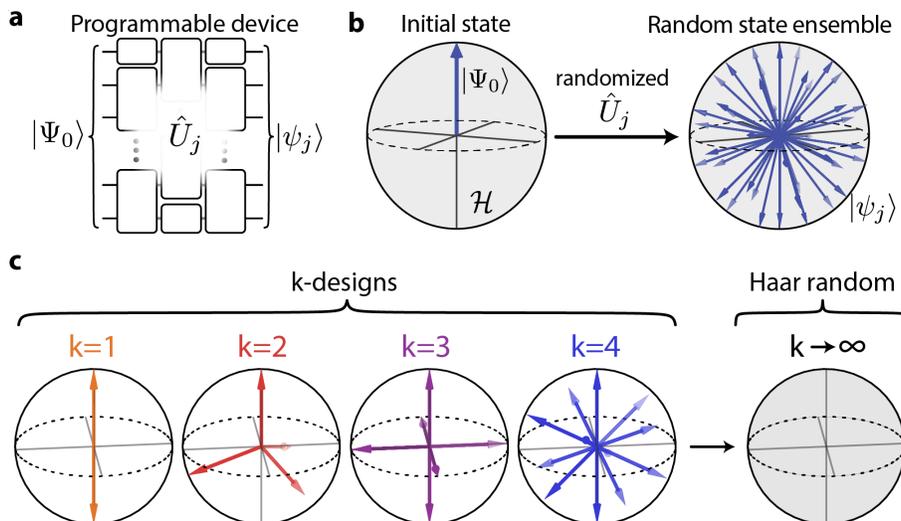


Figure 5.1: **Random pure state ensembles from Hamiltonian dynamics.** **a**, A thought experiment, consisting of a programmable device producing arbitrary quantum states $|\psi_j\rangle$ through unitary operations \hat{U}_j , where j enumerates over different program setting. **b**, Repeatedly applying explicitly randomized unitary evolution to an initial state $|\Psi_0\rangle$ produces an ensemble of pure quantum states $|\psi_j\rangle$ (blue arrows) which is distributed near-uniformly over the Hilbert space, \mathcal{H} (grey sphere), a *random state ensemble*. **c**, Practical quantum devices cannot sample from perfectly random state distributions, but only from low-order approximations. These ensembles are known as k -designs, and approach the fully Haar-random ensemble.

– is homogeneously distributed over the output Hilbert space, it is termed a Haar-random (or uniform) state ensemble [285], which we denote as \mathcal{E}_∞ (Fig. 5.1b). More formally, an ensemble of states is considered Haar-random if any unitary rotation of the ensemble leaves it invariant. It is easiest to think of the one-qubit Bloch sphere picture, and realize that this statement essentially just means the ensemble of states covers the entire Bloch sphere surface.

Practically, approximations to Haar-random state ensembles are conventionally generated by certain quantum devices requiring explicit classical randomization, in the sense that output states $|\psi_j\rangle$ are produced by randomly chosen unitary evolution operators \hat{U}_j . Examples include random unitary circuits [267, 282], where each configuration j is realized by a random choice of single and two-qubit gates, and stochastic evolution with a dynamically changing Hamiltonian [284, 286], $\hat{H}_j(t)$. However, it is currently unknown how to generate such random ensembles from the simplest form of quantum evolution, that governed by a fixed, time-independent Hamiltonian \hat{H} which is not explicitly randomized, as is the case for dynamics of closed and unperturbed quantum systems. In the next section by considering

pure state ensembles generated during partial measurement of a larger quantum system, we show such random ensembles *do* in fact emerge under such conditions.

It is important to emphasize this conceptual point. Normally in quantum mechanics, we start with an initial state, apply some deterministic evolution, and end with a fixed final state. Now, the setting we are considering is different. We are starting with a fixed initial state, then we are *sampling* different forms of unitary evolution (each with a probability p_j), which respectively result in different final states $|\psi_j\rangle$. The object of our study is not any individual $|\psi_j\rangle$, but instead the *ensemble*, \mathcal{E} , over all of them, which we can for instance denote as $\mathcal{E} \equiv \{[p_j, |\psi_j\rangle]\}_j$, where we include the p_j explicitly (sometimes we will omit them for notational simplicity). Importantly, we can define a given ensemble however we like, for instance we could define an ensemble of just a single qubit as $\mathcal{E}_1 = \{[0.5, |0\rangle], [0.5, |1\rangle]\}$, where the two states both have equal probability of 0.5 (left-most panel of Fig. 5.1c). Intuitively we can understand this ensemble is ‘less complex’ than an ensemble over all possible states on the Bloch sphere, a concept we will formalize below. Conceptually we can then think of our quantum device as *producing* the ensemble \mathcal{E} , and then randomly *sampling* a given $|\psi_j\rangle$ from it. While this may seem a subtle distinction, it will be a helpful mindset in future discussions, and also better clarifies what we mean by ‘preparing random state ensembles.’

It is *extremely* important to understand the distinctions between: 1) an ensemble of pure states, 2) a superposition of pure states, 3) a mixed state. First, there is no notion of superposition in our definition of an ensemble. We could have composed our ensemble of any possible pure state, including those which are themselves superpositions in a given basis, but the ensemble itself is just a *collection* of these pure states, not a superposition of them. Then, even more importantly, a state ensemble is *not* a mixed state. You can certainly define a mixed state *from* a state ensemble just via the sum $\hat{\rho} = \sum_j p_j |\psi_j\rangle\langle\psi_j|$, but the state ensemble contains far more information than the mixed state (even if the states $|\psi_j\rangle$ are orthogonal!). Indeed, as we will define further below, the mixed state can be thought of as the ‘mean’ of the ensemble, whereas the ensemble itself encodes higher-order information (analogous to higher-order quantities like variance, etc.). This is perhaps the most important fundamental concept to understand, so it bears repeating, and rebolding:

A state ensemble is not a mixed state.

With those important conceptual ideas presented, it is important to mention that notions of Haar-randomness have been studied extensively theoretically, and there are many very

formal ways in which to contextualize it via the lenses of Lie algebras, group theory, etc. For the most part, we will be forgoing such formalities here, and will take Haar-random to simply mean “random over the Hilbert space.” Readers interested in delving more into the theoretical subtleties are encouraged to peruse Refs. [285, 287], amongst others.

It is worth discussing sampling states from an ensemble, and specifically ask how does one sample a pure state from a Haar-random ensemble? Well, the Haar-measure is essentially the symmetry of a circle: any unitary operation, i.e. rotation, leaves it invariant. Thus we want to sample uniformly on the perimeter of a (high-dimensional) circle, but our basis (the elements of our state vector) are effectively Cartesian. This problem is well known mathematically, and an efficient solution is known as the Box-Muller transform [288]. As the simplest example, consider a two dimensional space. We sample both x and y (our Cartesian coordinates) from a standard normal distribution (i.e. with mean 0 and standard deviation 1). The resulting multi-dimensional Gaussian factorizes in rotational and radial components, so when we normalize the resulting vector, we find that it is uniformly sampled on the perimeter of the unit circle. The same algorithm applies efficiently for generating high-dimensional spherically-random (i.e. Haar-random) vectors. Explicitly, a D -dimensional Haar-random state is a complex Gaussian-random vector. In other words, generate a D -dimensional vector, where every element is a Gaussian-random variable with mean 0 and standard deviation 1. Add this vector to another Gaussian-random vector multiplied by i . Normalize, and the result is Haar-random. We further describe how to generate Haar-random unitaries in Appendix B. Also, we note that this connection to Gaussian random variables is explicitly used in many of the proofs in Ref. [285].

In the rest of this section, we briefly summarize the key observables and features of random state ensembles. Then, in the next section, we describe how we discover that random state ensemble behavior can emerge from unitary, deterministic quantum evolution.

5.1.1 Observables of random ensembles

With the concept of Haar-random state ensembles in mind, it is important to define their observable characteristics, so we know what experimental signatures to study. Two observables stand out in their universality and usefulness: moments of Haar-random observables, and the Porter-Thomas [289] distribution.

Moments The first observable of interest are the moments of the state ensemble. Staying first on the level of observables, moments are the fluctuations of those observables across

the ensemble. For instance, consider again our ensemble on a single qubit, defined above, $\mathcal{E}_1 = \{[0.5, |0\rangle], [0.5, |1\rangle]\}$. Say we wanted to calculate the mean probability to sample the ensemble, and find a qubit in the state $|1\rangle$. Obviously, this is just $m_1^{(1)} = 0.5(0) + 0.5(1) = 1/2$. This mean is the first moment of the ensemble; we indicate the order of the moment in the subscript on m , while we indicate that it is taken over \mathcal{E}_1 via the superscript.

In the same way as we talk about higher-order moments of probability distributions (like the mean, variance, skewness, kurtosis, etc.)¹, we can define higher order moments of the state ensemble. For instance, what is the second moment of \mathcal{E}_1 ? It is again simply calculated as $m_2^{(1)} = 0.5(0)^2 + 0.5(1)^2 = 1/2$, and in fact every k -th moment will also be $m_k^{(1)} = 1/2$. To be concrete, what we are studying is: if you sampled a random pure state many times from the ensemble \mathcal{E}_1 , and each time you measured if the qubit was in the up state for that pure state, what would the variance² of those measurements be over the set of all sampled pure states? We could ask the same question for any higher-order moment, where we are essentially asking: how *fluctuating* are observables of this ensemble?

Having understood this basic idea, let us ask what the moments of the Haar-random ensemble, \mathcal{E}_∞ , are. We shall not go into the full derivation here [285], but at least for our simple one-qubit manifold we have so far considered it is a simple spherical integral. For instance, in a Bloch sphere picture the probability to be in the state $|1\rangle$ is $\cos(\theta/2)^2$; if we then want the k -th order fluctuations of this quantity, we simply compute the normalized spherical integral explicitly

$$\frac{1}{4\pi} \int_0^\pi \int_0^{2\pi} \cos(\theta/2)^{2k} \sin(\theta) d\phi d\theta = \frac{1}{k+1}. \quad (5.1)$$

To generalize this problem, imagine that we are in a D -dimensional Hilbert space, for instance the tensor product space of several qubits. For a given pure vector, $|\psi\rangle$, say we are interested in the probability to measure a given bitstring, z , which we denote as $r_\psi(z) = |\langle z|\psi\rangle|^2$. To compute this more general case we use Gaussian integrals [285] to find the k -th moment over \mathcal{E}_∞ is

$$m_k^{(\infty)} \equiv \mathbb{E}_{\psi \in \mathcal{E}_\infty} [r_\psi(z)^k] = \frac{k!}{D(D+1) \cdots (D+k-1)} = \frac{1}{\binom{D+k-1}{k}}. \quad (5.2)$$

¹Note that variance is a *central moment* (i.e. one in which the mean has been subtracted off), whereas primarily here we will be discussing *raw moments*. For instance, for a random variable X , the k -th raw moment is just $\mathbb{E}[X^k]$. Normally people are more familiar and comfortable with central moments, which is why we list them here as an example. Of course, skewness and kurtosis are not even central moments, they are standardized moments, but we will not get into that here.

²Again, we are only using the word variance here colloquially. We actually are studying the raw moments, not the central moments.

One can then immediately verify this matches with our simplified spherical integration for $D = 2$ (i.e. for a single qubit). The most important thing to remember, conceptually, is that though this expression looks more complicated, it is simply a ‘spherical integral’ in a higher dimensional manifold, essentially.

There is one final thing to mention about the moments worth mentioning, which we pose as a temporary mystery to motivate our upcoming discussion.

Let us compare the first few moments of \mathcal{E}_1 (the ensemble of antipodal states on the Bloch sphere) and \mathcal{E}_∞ (the ensemble over all states on the surface of the Bloch sphere). Before we found $m_k^{(1)} = 1/2$, but $m_k^{(1)} = 1/(k + 1)$, but these expressions only match for $k = 1$. Now, let us somewhat *sans*-motivation propose a different ensemble of states \mathcal{E}_3 , which is composed of the six cardinal states with equal probabilities, i.e. $\mathcal{E}_3 \equiv \{|1\rangle, |0\rangle, |+\rangle, |-\rangle, |+i\rangle, |-i\rangle\}$ (see the third panel from the left in Fig. 5.1c). What are the first few moments of this ensemble?

$$m_k^{(3)} = \frac{1}{6}(0) + \frac{4}{6}\left(\frac{1}{2}\right)^k + \frac{1}{6} = \frac{1 + 2^{2-k}}{6} \quad (5.3)$$

$$m_1^{(3)} = \frac{1}{2}, \quad m_2^{(3)} = \frac{1}{3}, \quad m_3^{(3)} = \frac{1}{4}, \quad m_4^{(3)} = \frac{5}{24} \dots \quad (5.4)$$

Thus, we see the moments of \mathcal{E}_3 match with the moments of \mathcal{E}_∞ up to the third moment¹, but past this point they diverge. In a moment, we will describe why this occurs, how we judiciously chose the element of \mathcal{E}_3 to make this statement true, and why this is important for our purposes.

The Porter-Thomas distribution Above we studied moments, or fluctuations of ensembles of states. However, the presence of moments implies there must be a distribution of sorts, which for reasons discussed below we will call the *Porter-Thomas* (PT) distribution [289]. What is this distribution?

Remember, previously we calculated moments of the quantity $r_\psi(z)$, which is itself a probability. Thus, the distribution we are interested in is a *probability-of-probabilities* (PoP) distribution, which is sometimes called the ‘*histogram-of-histograms*’ or the ‘*fingerprint*’ amongst many other names in the statistics community [290].

This concept can seem unintuitive at first, so first we return to the simple example of \mathcal{E}_1 , where we studied, $r_\psi(1)$, the probability to measure the qubit in the up state, $|1\rangle$. There are two possible states we can sample from \mathcal{E}_1 , with corresponding $r_{|0\rangle}(1) = 0$ and $r_{|1\rangle}(1) = 1$. Thus, the PoP distribution in this case is a set of two delta functions, at 0 and 1.

¹Which might suspiciously motivate our decision to give \mathcal{E}_3 a subscript of 3...

More generally, we could have a system of N qubits, and could study $r_\psi(z)$ for any given bitstring z . For a given state ensemble, \mathcal{E} , composed of J total pure states $|\psi_j\rangle$ with equal probability then we can enumerate all the $r_{\psi_1}(z), r_{\psi_2}(z), \dots, r_{\psi_J}(z)$, and the PoP distribution will then be the sum of delta functions at each of these r values. More concretely, for a given ensemble, we can write the PoP as

$$P_{\mathcal{E}}(p) \equiv \mathbb{E}_{\psi \in \mathcal{E}}[\delta(r_\psi(z) - p)]. \quad (5.5)$$

Specializing to the case of the Haar-random ensemble, we can then easily write the moments of this distribution as $m_k^{(\infty)} \equiv \int_0^1 P_\infty(p)p^k = 1/\binom{D+k-1}{k}$, where we have used Eq. (5.2). We can then recognize¹ the result as the moments of a Beta distribution, specifically

$$P_\infty(p) = (D-1)(1-p)^{D-2}. \quad (5.6)$$

Importantly, in the limit of $D \rightarrow \infty$, this distribution becomes [291] the exponential distribution, $P_\infty(p) \approx \exp(-Dp)$.

The Porter-Thomas distribution is a well-studied object, particularly in the realm of *random matrix theory* [292–294]. It is a prototypical emergent observable of quantum chaotic systems [295]. However, it also appears in much more ‘pedestrian settings’: for instance, if you shine a laser on rough surface, you will see a ‘speckle pattern,’ a seemingly random pattern of constructive and destructive quantum interferences, for which it turns out the intensity distribution is also a Porter-Thomas distribution [296]. Basically, the Porter-Thomas distribution (which, again, is just a fancy name for an exponential distribution) shows up whenever quantum systems are randomized, not in the sense of their measurement, but in the sense of their *evolution*². A full description of the Porter-Thomas distribution’s properties and appearances throughout physics is certainly beyond the scope of this thesis: our only real claim about it here is that it is the distribution which describes observables drawn from a Haar-random state ensemble.

However, there is something else interesting to note, which is not immediately obvious from our derivation, and which bears stating. Above, we claimed that the probability to measure a single, fixed bitstring, z , from the Haar-random ensemble of states was distributed as the Porter-Thomas distribution. Now, imagine we take D to be large (for a tensor product

¹We will not be deriving this whole process here. However, the basic schematic is that one knows all the moments, m_k , one can write down the *moment-generating function*, $M(t) = 1 + m_1 t + m_2 \frac{t^2}{2!} + \dots + m_n \frac{t^n}{n!} + \dots$. Then, one can form the *characteristic function*, $\psi(t) = M(it)$, and finally one takes the Fourier transform to arrive at the probability density function, which we provide above.

²As a slight spoiler for what is about to come, our central discovery in the next section is that the entanglement – in combination with the randomness of Born’s rule – *also* leads to fundamentally quantum description of random quantum evolution for generic many-body states.

space of qubits), and we sample a single pure state, $|\Psi\rangle$ from the Haar-random ensemble, and we consider the probability to measure every single possible bitstring from $|\Psi\rangle$, i.e. $r_\Psi(z_0), r_\Psi(z_1), \dots, r_\Psi(z_{2^N})$. If we now form a histogram of all these r , we will find that *it too* follows the Porter-Thomas distribution. In other words, the probability distribution over bitstrings for a single typical Haar-random state is the same as the probability distribution for a single bitstring over the entire Haar ensemble. This notion of ‘typicality’ is remarkable, and will be important in many future discussions. We do not present a full derivation of it here, but we do note that it is derived, very pedagogically and thoroughly, in Ref. [297].

Having defined the Porter-Thomas distribution and the moments of the Haar-random ensemble, we now turn to better define what it means for an ensemble to be complex, and how complex can real experiments be.

5.1.2 Quantum state k -designs

It is important to highlight that in practice most forms of quantum evolution are incapable of sampling truly uniformly from the available Hilbert space, and thus are not strictly speaking Haar-random. Thus, it is important to develop a metric of how random they in fact are. For this reason, we will often study so-called *quantum state k -designs* [298], which are low-order approximations of the Haar-random ensemble.

Quantum state k -designs are pure state ensembles that become progressively more complex with larger k ; for example, for the single-qubit case shown in Fig. 5.6a, state k -designs appear as increasingly uniform distributions on the Bloch sphere, and realize the uniform ensemble in the infinite k limit. A state k -design reproduces observables up to order k as if they were sampled from the uniform ensemble, e.g., the analytic result of for the moments discussed in Section 5.1.1 can be predicted from a state k -design. That is why \mathcal{E}_1 (which is a 1-design) was able to match the moments of the Haar-random ensemble only up to first order, but \mathcal{E}_3 (which is a 3-design) was able to do so up to third order. A vanishing distance between the state k -design and a given ensemble implies that the ensemble and the Haar-random ensemble are statistically indistinguishable for any observables up to order k .

It is important to highlight that this is not just a statement about observables of the random state ensemble, but about the ensemble itself. If one only has access to the average of the random state ensemble (which would be a mixed state), one would not be able to calculate moments beyond first order (e.g. the average observable obtained from the ensemble). We can find a generalized description however via the k -th moments of the ensemble itself. Consider if we have an ensemble \mathcal{E} over states $\{|\psi_j\rangle\}$ and associated sampling probabilities

p_j . Then in that case, we construct a higher-order version of a ‘mixed state’ via

$$\rho_{\mathcal{E}}^{(k)}(t) = \sum_j p_j (|\psi^j\rangle\langle\psi^j|)^{\otimes k}, \quad (5.7)$$

where k is the order of the ensemble. Note that, as discussed previously, the normal mixed state is simply the $k = 1$ version of this generalization, further reiterating our declaration that

A state ensemble is not a mixed state.

This distinction is made very apparent here, because in order to describe an observable of order k from the ensemble (such as k -th order moments of observables drawn from the ensemble), one formally needs to employ the higher-order mixed state generalization

$$m_k = \text{Tr}[\rho_{\mathcal{E}}^{(k)}(|z\rangle\langle z|)^{\otimes k}] \quad (5.8)$$

where Tr is the trace operation. If one only has access to the mixed state, $k = 1$, then one can only describe observables up to order $k = 1$, like the average value measured from an ensemble.

So what are the $\rho_{\mathcal{E}}^{(k)}$ for the Haar-random ensemble, which we will denote $\rho_{\infty}^{(k)}$? Well, again we start with our simple example of a one-qubit system. Here, we can explicitly write the state vector in spherical coordinates as

$$|\psi\rangle = \begin{pmatrix} \cos(\theta/2) \\ \sin(\theta/2)e^{i\phi} \end{pmatrix}. \quad (5.9)$$

Correspondingly, the state $|\psi\rangle^{\otimes 2} \equiv |\psi\rangle \otimes |\psi\rangle$ is given by

$$|\psi\rangle^{\otimes 2} = \begin{pmatrix} \cos(\theta/2)^2 \\ \cos(\theta/2)\sin(\theta/2)e^{i\phi} \\ \cos(\theta/2)\sin(\theta/2)e^{i\phi} \\ \sin(\theta/2)^2 e^{2i\phi} \end{pmatrix}. \quad (5.10)$$

We can then explicitly integrate $|\psi\rangle^{\otimes 2}$ over the surface of the Bloch sphere, finding

$$\frac{1}{4\pi} \int_0^\pi \int_0^{2\pi} (|\psi\rangle\langle\psi|)^{\otimes 2} \sin(\theta) d\phi d\theta = \frac{1}{6} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} = \frac{1}{6}(\hat{I} + \hat{S}), \quad (5.11)$$

where \hat{I} is the diagonal identity operator, and \hat{S} is the SWAP operator. One could repeat this calculation for $|\psi\rangle^{\otimes 3}$, in which case you would find

$$\begin{aligned} \frac{1}{4\pi} \int_0^\pi \int_0^{2\pi} (|\psi\rangle\langle\psi|)^{\otimes 2} \sin(\theta) d\phi d\theta &= \frac{1}{24} \begin{pmatrix} 6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 2 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & 2 & 0 \\ 0 & 2 & 2 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & 2 & 0 \\ 0 & 0 & 0 & 2 & 0 & 2 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix} & (5.12) \\ &= \frac{1}{24} (\hat{S}_{123} + \hat{S}_{132} + \hat{S}_{213} + \hat{S}_{231} + \hat{S}_{312} + \hat{S}_{321}), & (5.13) \end{aligned}$$

where the \hat{S} are permutation operators, e.g. \hat{S}_{321} is the operator which swaps the 1st and 3rd copies of the state vector in the Kronecker product, leaving the 2nd copy fixed:

$$\hat{S}_{321} = \begin{pmatrix} \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 \\ 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 \\ 0 & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \mathbf{1} & 0 & 0 \\ 0 & 0 & 0 & \mathbf{1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{1} \end{pmatrix}, \quad \vec{b} = \begin{pmatrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{pmatrix}, \quad (5.14)$$

where we have bolded the matrix entries to improve visibility, and have further written the binary representation of the basis vector, \vec{b} to make more clear the correspondence mapping.

One could continue in this way infinitely and in general would find

$$\rho_\infty^{(k)} = \frac{1}{(k+1)!} \sum_{\sigma \in \mathcal{P}_k} \hat{S}_\sigma \quad (5.15)$$

where \mathcal{P}_k is the group of permutations of k elements. Fancier math [285] can do the same calculation for arbitrary dimensional Hilbert space, yielding the same basic answer, but with a different normalization constant

$$\rho_\infty^{(k)} = \frac{\sum_{\sigma \in \mathcal{P}_k} \hat{S}_\sigma}{D(D+1) \cdots (D+k-1)}. \quad (5.16)$$

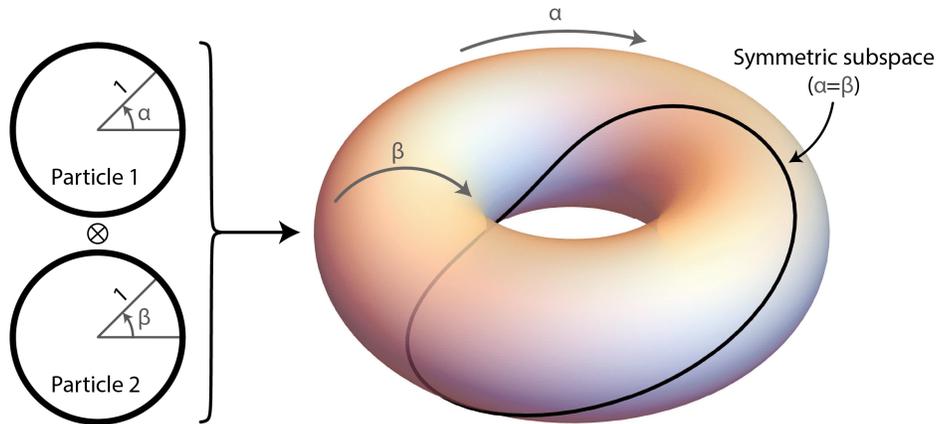


Figure 5.2: **The symmetric subspace.** When we integrate over k copies of Hilbert space, for instance to calculate the k -th moments of the Haar-random ensemble, the integration takes place in a special manifold of Hilbert space: the symmetric subspace [285]. Here we visualize this for an exclusively real qubit Hilbert space, i.e. one in which the Hilbert space is a circle rather than the Bloch sphere, for which the tensor product space is a torus. Note that generically for complex, multi-qubit Hilbert space, the space topology is a hyper-torus. The symmetric subspace (black line) is the unique manifold with no symmetries, except the permutation symmetries between the k copies. You can convince yourself visually that averaging over the symmetric subspace (the black line) *does not* give the same mean as averaging over the entire Hilbert space (the surface of the torus). The average of the torus is the identity operator (the origin), whereas the average of the symmetric subspace is offset by permutation operators, see Eq. (5.16). (In this case the average is offset far more drastically because we are averaging over real Hilbert spaces, see Ref. [285].)

So what does this mean?

What we are finding is that the k -th order generalization of a mixed state of the Haar-random ensemble is essentially the sum over all possible permutations of k copies of the Hilbert space. Hopefully this result is intuitive: when we integrate over the set of all Haar-random states, we are integrating uniformly over the Hilbert space. This process averages over all possible symmetries leaving only the identity (i.e. the maximally mixed state). However, there is still a symmetry which is not averaged over: the permutation symmetry between the different copies of Hilbert space. This residual symmetry means that averages over k -th moments of the Haar-ensemble have additional structure beyond the maximally mixed state. Thus, we work in the so-called *symmetric subspace* [285], the sub-manifold of the k -copied Hilbert space with unbroken permutation symmetry (Fig. 5.2).

This property that averages over k copies of the Haar-random ensemble can be reduced to sums of permutation operators can greatly simplify many calculations, some of which we

outline below.

5.1.3 Why does randomness matter?

Why is it important to know more than the mean of the ensemble? Because understanding the fluctuations of an observation is just as important as the observation itself, not only for bounding the amount by which observables can vary, but for practical purposes. This is because of what is sometimes referred to as the *swap tricks* [287], that for the two-body SWAP operator \hat{S}_{21} , any operators \hat{A} and \hat{B} , and any states $|\phi\rangle$ and $|\psi\rangle$ we can write

$$\hat{S}_{21}(|\psi\rangle \otimes |\phi\rangle) = |\phi\rangle \otimes |\psi\rangle \quad (5.17)$$

$$\text{Tr}((\hat{A} \otimes \hat{B})\hat{S}_{21}) = \text{Tr}(\hat{A}\hat{B}) \quad (5.18)$$

$$\text{Tr}_2((\hat{A} \otimes \hat{B})\hat{S}_{21}) = \hat{A}\hat{B}, \quad (5.19)$$

where Tr_2 is the partial trace over the second of the tensor product spaces. The second two of these properties, the action of the SWAP operator on tensor products of operators, is particularly useful, as it allows us to contort a tensor network representation of a given quantum operation into a more easily computable form [30, 287, 299]. This for instance is used in the classical shadows formalism [300] to both estimate arbitrary observables of a quantum system from a few randomized measurements, and also to bound the confidence of those estimations; see for instance Ref. [287]. Further, it can be used calculate the purity efficiently of a multipartite states averaged over random states, again see Ref. [287].

There are of course also higher-order analogs of the swap trick. For instance if we consider cyclic permutation operators, $\hat{C}_d^{(k)}$, which are permutations which keep the order of the k indices the same but just rotate them by d sites e.g. $\hat{C}_2^{(4)} = \hat{S}_{3412}$, then we can write the generalized swap trick

$$\text{Tr}((\hat{A}_1 \otimes \cdots \otimes \hat{A}_k)\hat{C}_1^{(k)}) = \text{Tr}(\hat{A}_1 \cdots \hat{A}_k). \quad (5.20)$$

One of the most important use-cases of the swap tricks is for benchmarking [58, 272–277], where we use swap tricks to change measurements of a quantum overlap fidelity between two quantum states into a cross-correlation between observables of those states. We will discuss this application at length in Chapter 6, but here present a qualitative justification. If we have a 1-design of a single qubit, \mathcal{E}_1 , composed of the antipodal states $|0\rangle$ and $|1\rangle$. As discussed, \mathcal{E}_1 only matches observables from the Haar-ensemble up to first order, and so is essentially just described by the maximally mixed state, $\hat{\rho} = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|$. Imagine that we now subject \mathcal{E}_1 to a depolarization channel, transforming each pure state itself to

the maximally mixed state, i.e. we define a generalized notion of a *mixed state ensemble*, which is defined as the action of some noisy channel on all the states of \mathcal{E}_1 . Well, for \mathcal{E}_1 there is no substantive difference after depolarization in terms of the correspondence to the Haar random ensemble: it is still only as ‘random’ as the maximally mixed state.

Now however, consider if instead of \mathcal{E}_1 , we started with \mathcal{E}_3 composed of the six cardinal states. Imagine this ensemble is subjected to the same depolarization channel which maps all inputs to the maximally mixed state. If we measured the second or third moment before depolarization we would find they matched the Haar-random expectation, but after depolarization they would not. In other words, we are using the fact that we can predict the moments of Haar-random ensembles to tell how far a given ensemble of (potentially mixed) states is from the maximally mixed state. The reader can hopefully then see that by studying how close observables match their Haar-random predictions, we can develop methods (which we can quantitatively formalize using swap tricks) of benchmarking the fidelity of a quantum operation under a noisy channel: we shall return to this idea extensively in Chapter 6. Later, in Chapter 7, we will show how by studying higher-order fluctuations of the ensemble produced by a noisy channel, we can actually learn the noise processes affecting our quantum experiment.

5.1.4 Relation to spherical designs

While we have tried to present intuitive, low-order explanations of Haar-randomness and k -designs, we understand these concepts may feel exotic still. Thus, here we finally make a concrete connection between this relatively abstract discussion of ensembles and k -designs, and remind the reader that ultimately, we are essentially asking how uniformly is a set of states spread over a given Hilbert space. Even more simply, how uniformly is a set of vectors spread over some geometric space. It turns out that the entire notion of k -designs can be captured in this simpler geometric picture, and in fact the very notion of k -designs has its origins in asking how faithfully an integral over a sphere can be reproduced through sampling only a finite set of points. Indeed, even the name of ‘ k -designs’ comes from the earlier study of ‘spherical-designs’ [301].

Concretely, we define a series of polynomials of three variables, the first few of which we will explicitly write out as

$$P_0(x, y, z) = a$$

$$P_1(x, y, z) = b_1x + b_2y + b_3z$$

$$P_2(x, y, z) = c_1xy + c_2yz + c_3zx + c_4x^2 + c_5y^2 + c_6z^2$$

$$P_3(x, y, z) = d_1x^2y + d_2y^2z + d_3z^2x + d_4x^2z + d_5y^2x + d_6z^2y + d_7x^3 + d_8y^3 + d_9z^3 + d_{10}xyz,$$

and where for more notation we define $\hat{P}_n = \sum_{i=0}^n P_i(x, y, z)$. Now consider we want to integrate such functions, say \hat{P}_1 , on \mathbb{S}^2 , the surface of a sphere. To do so, we could perform the full spherical integration, and in doing so calculate \hat{P}_1 for every possible point on the sphere. Alternatively, however, we could simply recognize that due to symmetry and the simplicity of \hat{P}_1 we know the answer is a (where we have normalized by the sphere surface area). In other words, what we have done is recognize that due to symmetry, integration over the sphere's surface is the same as averaging over just two antipodal points, i.e.

$$\mathbb{E}_{\{r_1\}}(\hat{P}_1) = a = \int_S \hat{P}_1(x, y, z) dx dy dz, \quad (5.21)$$

where S represents the surface of the sphere (and again we are normalizing the sphere surface area), $\{r\}$ is the set of points we are averaging over, \mathbb{E} is an expectation value, and $\{r_1\} = \{(0, 0, 1), (0, 0, -1)\}$.

What about, however, if we tried integrating a higher order polynomial, like \hat{P}_2 with the same method? Ostensibly the same spherical symmetry holds, so we could try the same trick as before, but we would find that

$$\mathbb{E}_{\{r_1\}}(\hat{P}_2) = a + c_5 \neq a + \frac{1}{3}(c_4 + c_5 + c_6) = \int_S \hat{P}_2(x, y, z) dx dy dz. \quad (5.22)$$

Clearly, something has gone awry, however, if instead of averaging over $\{r_1\}$ we instead averaged over $\{r_2\}$, the vertices of a regular tetrahedron embedded in the sphere, we would find $\mathbb{E}_{\{r_2\}}(\hat{P}_2) = a + \frac{1}{3}(c_4 + c_5 + c_6)$ as desired. To belabor the point but a moment longer, if we tried integrating \hat{P}_3 , we would now find $\{r_2\}$ insufficient, and would have to average over $\{r_3\}$, the vertices of a regular octahedron.

What we are essentially discovering is that even for complex functions, we can forgo integration over the entire space and instead calculate our function at a few judiciously chosen points (the number and placement of which depends on how complex the function is). Notably, using a higher-order choice of points will also yield the correct answer for lower-order polynomials, i.e. averaging over $\{r_3\}$ would yield the correct answer for integrating \hat{P}_1 or \hat{P}_2 .

The punchline of all of this is that the same exact intuition that applies to integrating polynomials on a sphere applies equally to calculating moments of observables of the

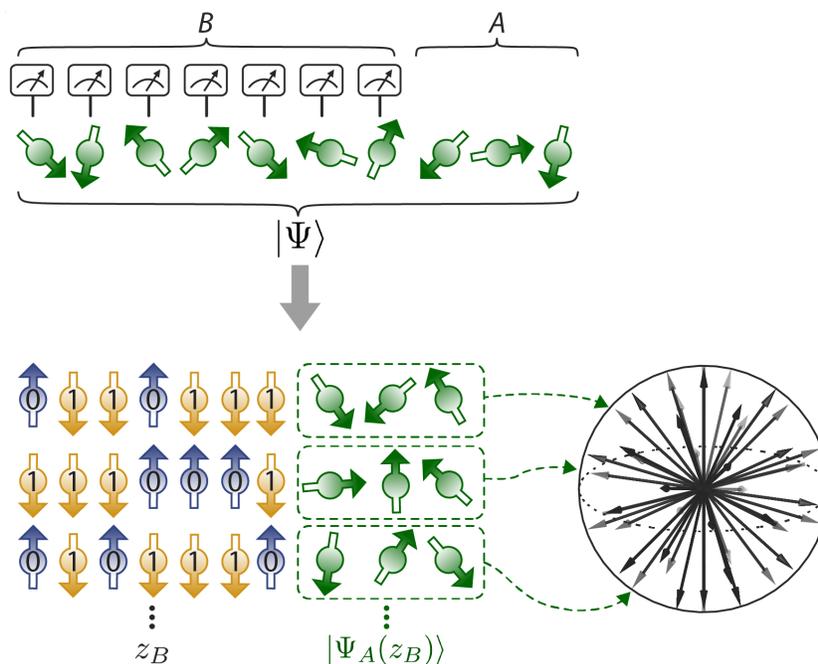


Figure 5.3: **Pictorial representation of the projected ensemble.** Partial measurement of a subset (B) of a quantum state projects those qubits into well-defined classical bitstrings, subject to the Born rule. However, entanglement between that subset and the rest of the unmeasured system (A) then ‘transfers’ the randomness of the Born rule onto the quantum states in A . The result is an ensemble of states, which we call the *projected ensemble*, tagged by classical bitstrings in B , which we find approach the Haar-random distribution for generic quantum evolution.

Haar-ensemble, which informally you can intuitively understand from the fact that the Haar-ensemble is a uniform cover of the sphere itself. For a precomputed list of spherical designs, which visually correspond to the set of states in k -designs on single qubit Bloch spheres (Fig. 5.1c), see Ref. [302].

5.2 THE PROJECTED ENSEMBLE

Having introduced the concept of random state ensembles, now we discuss our novel discovery of emergent randomness from unitary, time-independent Hamiltonian evolution. Remember, it is well known that stochastic, time-dependent evolution can produce (nearly) Haar-random states in exponential time [267, 282, 284, 286] - we are interested in the same can somehow be said for evolution without injecting classical randomness, instead relying only on the most fundamental features of quantum mechanics. We present a new perspective on the emergence of statistical behavior in chaotic quantum many-body systems: instead of imagining that a physical state is sampled from a random ensemble, we use a single

many-body wavefunction to *generate* an ensemble of pure states on a subsystem (Fig. 5.3).

To study the emergence of random state ensembles, we generically consider Hamiltonian evolution that produces a global quantum state $|\psi\rangle$, which we here suppose describes a set of qubits with basis states $|0\rangle$ and $|1\rangle$. We bipartition the state into two subsystems: a local system of interest, A , and its complement B (Fig. 5.4c). Explicitly keeping track of measurement results in B , which are bitstrings of the form e.g. $z_B = 100 \cdots 010$, provides a full description of the total system state as

$$|\psi\rangle = \sum_{z_B} \sqrt{p(z_B)} |\psi_A(z_B)\rangle \otimes |z_B\rangle, \quad (5.23)$$

where $p(z_B)$ is the probability of measuring a given bitstring in B , and $|\psi_A(z_B)\rangle$ is a pure quantum state in A conditioned on the measurement outcome in B . Thus, for each possible z_B , there is a well-defined pure state in A , the set of all of which is generally not orthogonal. Together these states, $|\psi_A(z_B)\rangle$, and their respective probabilities, $p(z_B)$, form what we term the *projected ensemble* [70] (Fig. 5.4b); similar concepts also enter the definition of localizable entanglement [303], and in the concept of conditional wavefunctions [304]. By tracking the time evolution of the projected ensemble through both the states and probabilities which compose it, we can probe for signatures of the ensemble approaching a Haar-random distribution (Fig. 5.4c).

We stress that this concept is distinct from typical studies of equilibration in quantum many-body systems. There, the central object of interest is the reduced density operator on A , $\hat{\rho}_A = \text{Tr}_B(\hat{\rho})$, found from tracing out B from the full density operator $\hat{\rho}$. The reduced density operator can be constructed by averaging over the projected ensemble states, $\hat{\rho}_A = \sum_{z_B} p(z_B) |\psi_A(z_B)\rangle \langle \psi_A(z_B)|$, but as such can only provide information on the mean of ensemble observables, and never on the actual ensemble distribution itself (see Section 5.1.2).

To elucidate the importance of this distinction and reveal the emergence of random statistical properties of the projected ensemble, we employ a Rydberg analog quantum simulator [34, 36, 41], which provides high fidelity preparation, evolution, and readout [34]. After a variable evolution time, we perform site-resolved readout in a fixed measurement basis, yielding experimentally measured bitstrings, z , which we bipartition into bitstrings z_A and z_B .

Hamiltonian parameters are chosen such that, after a short settling time, the *marginal* probability, $p(z_A)$, of measuring a given z_A (while ignoring the complementary z_B) agrees with the prediction from $\hat{\rho}_A$ being a maximally mixed state. In the language of quantum

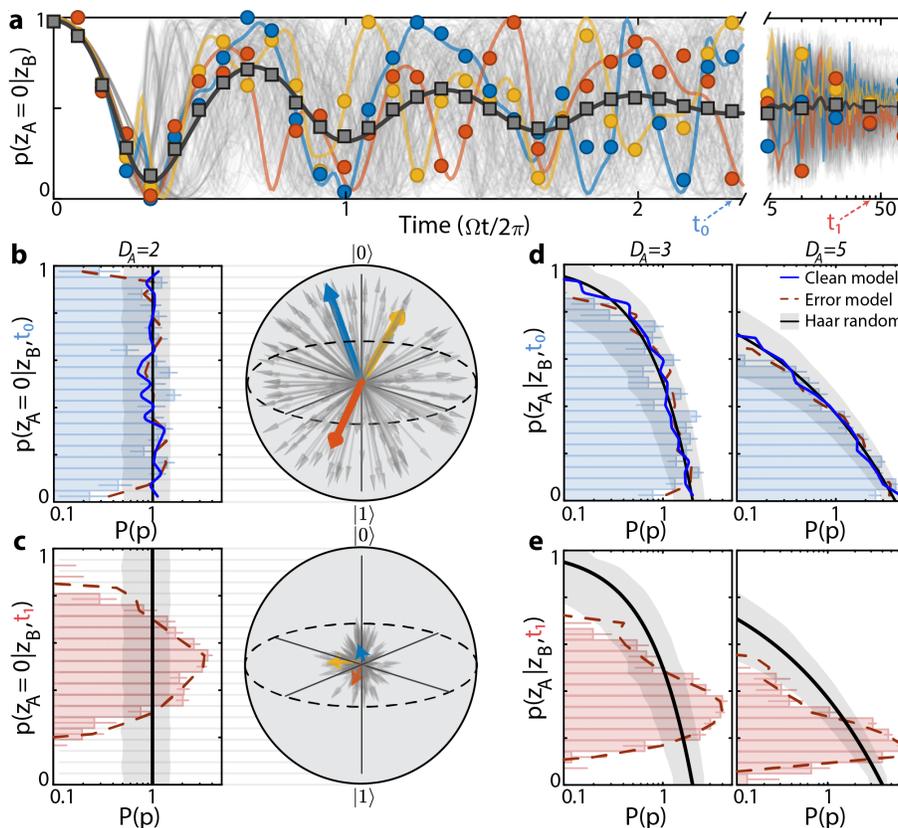


Figure 5.5: Experimental signatures of random pure state ensembles. **a**, We employ a ten-atom Rydberg quantum simulator to perform Hamiltonian evolution leading to quantum thermalization at infinite effective temperature (see Appendix C). For a single qubit in A , we plot the probabilities for finding a single qubit subsystem in state $|0\rangle$ as a function of time. Grey square markers indicate the *marginal* probabilities $p(z_A=0)$, which equilibrate to ~ 0.5 due to thermalization with B . In contrast, colored circle markers show *conditional* probabilities given a specific measured z_B in B , $p(z_A=0|z_B)$, which show large fluctuations even after the marginal probability reaches a steady state; these then diminish at late times due to extrinsic decoherence effects. Such conditional probabilities yield information about the projected ensemble as $p(z_A|z_B) = |\langle z_A|\psi_A(z_B)\rangle|^2$. Grey lines are simulated trajectories of $p(z_A=0|z_B)$ for all outcomes z_B , with a few highlighted to be compared with experimental data (color lines and markers). Decoherence sources are included for simulations after the axis break. **b**, Histograms, $P(p)$, of the probabilities $p(z_A=0|z_B)$ at intermediate ($\Omega t_0/2\pi = 2.3$) time. The experimental results are close to a flat distribution, consistent with a Haar-random ensemble, as visualized by the simulated distribution of projected states (right). **c**, However, at late ($\Omega t_1/2\pi = 38$) time, decoherence effects have concentrated probabilities around $1/D_A = 0.5$, consistent with the error model simulation showing the reduced lengths of single-qubit states (right). **d,e**, Similar agreement with predictions from random state ensembles is also seen for larger subsystem sizes of A with higher subsystem dimension, D_A . In **b-e**, black lines and grey bands are predictions and uncertainties (from finite sampling) of a D_A -dimensional uniform random ensemble; red dashed lines and blue solid lines are from simulations with and without decoherence

yields information of the reduced density operator, such conditional probabilities yield signatures of the projected ensemble, as $p(z_A|z_B) = |\langle z_A|\psi_A(z_B)\rangle|^2$. In Fig. 5.5a, we plot numerically simulated $p(z_A=0|z_B)$ in grey, with selected traces highlighted in color to be compared with their corresponding experimental data (circle markers). Importantly, conditional probabilities fluctuate chaotically with sensitive dependence on z_B , even when the marginal probability has reached a steady state. In experiments, fluctuations slowly damp out over time due to extrinsic decoherence effects from coupling to an *external* environment at very late time, but these decoherence effects do not appear to affect the late-time marginal probability (right panel, Fig. 5.5a).

For calculations involving conditional probabilities, we bipartition each bitstring z into subsystems A and B with bitstrings z_A and z_B respectively. When considering the statistics of conditional probabilities, we note that the blockade interaction [36] can reduce the dimensionality of the Hilbert space of subsystem A if the boundary qubits in B are in the Rydberg state. To isolate a set of conditional states having the same Hilbert space dimension, D_A , for a given choice of subsystem A and B , we only consider bitstrings z_A and z_B if the qubits in B bordering A are in the $|0\rangle$ state.

To analyze fluctuations, we bin conditional probabilities into a histogram $P(p)$ for a time when fluctuations are strong and decoherence effects are small (t_0 , Fig. 5.5b) as well as at very late time (t_1 , Fig. 5.5c) when decoherence dominates. At t_0 , the experimental $P(p)$ distribution is essentially flat, as predicted for a Haar-random ensemble (Eq. (5.6)), up to finite-sampling fluctuations and weak decoherence effects. We show projected states obtained from simulation (Bloch sphere in Fig. 5.5b), including decoherence, to illustrate how such a flat distribution is generated from a near-uniform ensemble of states. At very late time, t_1 , decoherence reduces the purity of projected states significantly, leading to $P(p)$ becoming concentrated around $1/D_A = 0.5$ (Fig. 5.5c). This highlights that the agreement between the experimental data and the random ensemble prediction in Fig. 5.5b,d is a coherent phenomenon of closed quantum system dynamics. We further validate this in Figs. 5.5d,e by plotting the $P(p)$ for A composed of 2 and 3 atoms, with corresponding Hilbert space dimensions of $D_A = 3$ and 5, respectively. Here, the prediction from the Haar-random distribution (Eq. (5.6)) is $P(p) = (D_A - 1)(1 - p)^{D_A - 2}$.

5.2.1 Moments of the projected ensemble

The convergence of the projected ensemble to a nearly Haar-random distribution can be temporally resolved by considering the moments of the distributions $P(p)$, where the k th moment is defined as $p^{(k)} = \sum_p p^k P(p)$ (Fig. 5.6a). Looking order-by-order, we find after

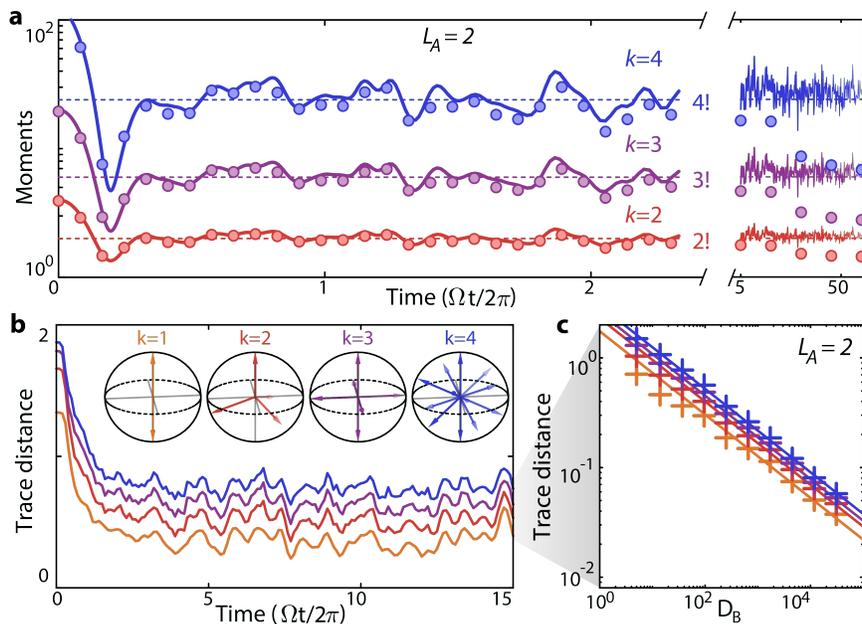


Figure 5.6: Development of emergent randomness. **a**, Rescaled second (red), third (purple), and fourth (blue) moments of the conditional probability distributions in Fig. 5.5b for subsystem of length $L_A = 2$. Experimental moments saturate to $\approx k!$, the expectation from the uniformly random ensemble (dotted lines) and consistent with numerical simulation (solid lines), before eventually decaying due to decoherence. **b**, Numerically computed trace distances as a function of time between the $L_A = 2$ projected ensemble and the four lowest order approximations to the uniform random ensemble, so called *quantum state k -designs*, for $k = 1, 2, 3, 4$ (inset). Distances for all k decrease initially before saturating due to finite system-size effects. If the trace distances up to order k vanish, the ensemble is as random as the k th design, and fluctuations of observables match up to order k , such as the k th moments in a. **c**, Late-time distances decrease as $\sim 1/\sqrt{D_B}$ (solid lines), the Hilbert space dimension of the effective bath, subsystem B .

rescaling by a factor of $D_A \cdots (D_A + k - 1)$, moments from both experiment and numerics quickly approach $k!$, the analytical result expected from a Haar-random ensemble (Eq. (5.2)). This is independent of the details of subsystem selection, whether A is chosen at the edge, center, or is even discontinuous (Fig. 5.7). At very late time, moments decrease, indicating sensitivity to decoherence effects (Fig. 5.6a, right). Further examples for different sizes of A are shown in Fig. 5.8.

Further, universal values are also found for two-point correlators, in agreement with their Haar-random predictions. In Fig. 5.9, we measure the square root of a weighted average of squared two-point correlators in the z -basis using the projected ensemble of a two-qubit

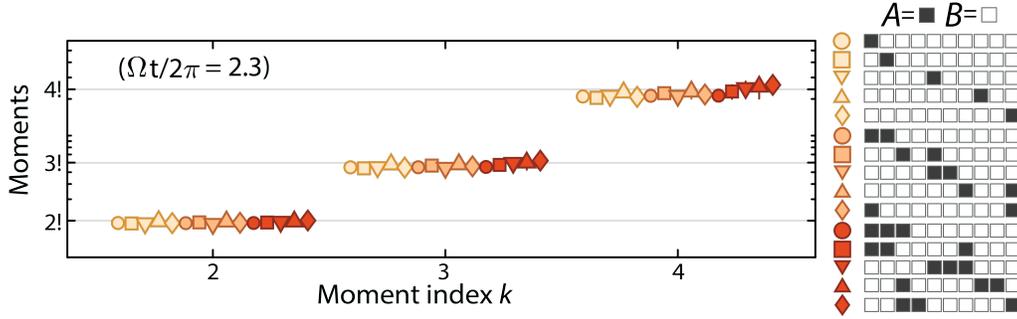


Figure 5.7: **Universality of moments of the projected ensemble.** k th moments of the conditional probability distributions in Fig. 5.5b,c, evaluated at late-time ($\Omega t/2\pi = 2.3$) and for a variety of choices of subsystems (see panel on the right); we find a universal convergence to $\approx k!$, independent of subsystem choice, suggesting that a subsystem's projected ensemble converges to the uniform random ensemble irrespective of the details of placement, or connectivity. Error bars are the standard deviation over temporal fluctuations in moments near the evaluated time, as shown in Fig. 5.6a.

subsystem, given as

$$\sigma_{\text{corr}}(z_1, z_2) = \sqrt{\sum_{z_B} p(z_B) C(z_1, z_2 | z_B)^2} \quad (5.24)$$

where

$$C(z_1, z_2 | z_B) = \langle \psi_A(z_B) | \hat{Z}_1 \hat{Z}_2 | \psi_A(z_B) \rangle - \langle \psi_A(z_B) | \hat{Z}_1 | \psi_A(z_B) \rangle \langle \psi_A(z_B) | \hat{Z}_2 | \psi_A(z_B) \rangle. \quad (5.25)$$

Here, \hat{Z}_i is the Pauli- Z operator at site i , $p(z_B)$ is the probability of observing a bitstring z_B from outer subsystem B , and the expectation values on the right-hand side are computed with the projected two-qubit state $|\psi_A(z_B)\rangle$. To compare the projected ensemble result with a Haar-random ensemble, we use Eq. (5.16) for $D_A = 4$ to compute the analytical value of such a correlator for the Haar-random ensemble:

$$\sqrt{\mathbb{E}_{\psi \in \text{Haar}(4)} \left[\left(\langle \psi | \hat{Z}_1 \hat{Z}_2 | \psi \rangle - \langle \psi | \hat{Z}_1 | \psi \rangle \langle \psi | \hat{Z}_2 | \psi \rangle \right)^2 \right]} = \sqrt{\frac{4}{35}}. \quad (5.26)$$

5.2.2 Comparison to quantum state k -designs

We study the trace distance between the projected ensemble, generated from error-free simulation, and successive k -designs (Fig. 5.6b), finding the distance decreases for all k th orders as a function of time, before saturating to a value exponentially small in the total

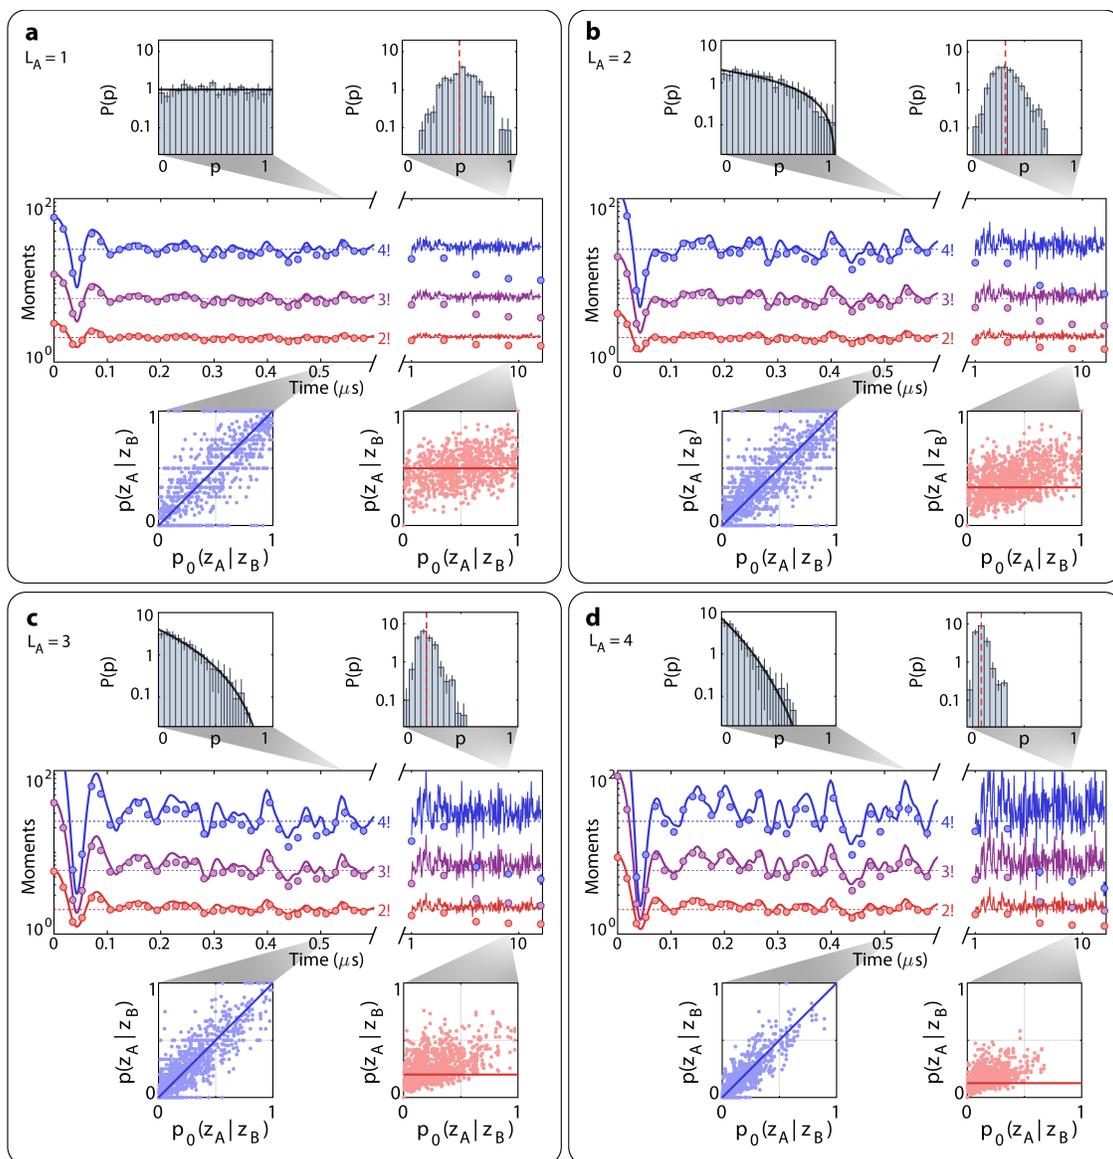


Figure 5.8: **Moments of the projected ensemble.** 2nd (red), 3rd (purple), and 4th (blue) moments of the projected ensemble for different choices of the size of A , ranging from $L_A = 1$ (top) to $L_A = 4$ (bottom). Data is taken from the same dataset as in Fig. 5.6

system size (Fig. 5.6c). Similar numerical results are found for the case of random unitary circuits and a Hamiltonian used in ion trap experiments (Fig. 5.10). In Ref. [70], we show that the formation of uniformly random, pure state ensembles in subsystems is a more universal phenomenon.

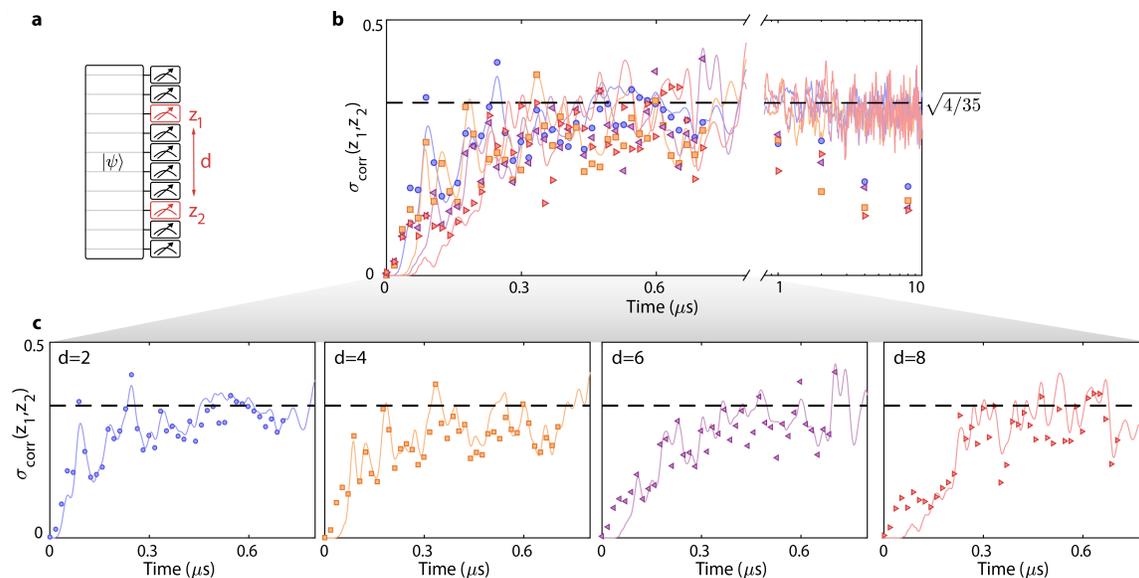


Figure 5.9: **Fluctuations of conditional two-point correlations in projected state ensembles:** **a**, Two-point correlations based on conditional projected states of two qubits separated at distance d . **b**, Fluctuations of two-point correlation functions of conditional probabilities, $\sigma_{\text{corr}}(z_1, z_2)$, characterized with a ten-qubit Rydberg-atom array simulator (Eq. (5.24)). Experimental (markers) and numerical (solid lines) results increase and saturate to near the theoretical prediction of $\sigma_{\text{corr}} = \sqrt{4/35}$ from the uniformly random ensemble. At late times $t > 1 \mu\text{s}$, the universal feature gradually vanishes due to accumulated errors and decoherence processes in the experiment (time scale is logarithmic after the break). Colors and markers used are: $d = 2$, blue circle; $d = 4$, orange square; $d = 6$, purple triangle; $d = 8$; red triangle. **c**, Individual plots for the four different pair separations shown in **b**.

5.2.3 Finite temperature projected ensembles

While all measurements thus far were conducted with the system thermalizing close to an infinite effective temperature condition, it is interesting to ask how the emergent randomness of the projected ensemble is changed in the finite effective temperature case (see Appendix C). We note that here we given only a very brief, and incomplete, discussion of this effect, and so for substantial further details we recommend the reader consult Ref. [71].

To study the finite temperature condition, we numerically simulate quench evolution under the Rydberg Hamiltonian with different choices of global detuning, and observe the effect on observables for an $N = 16$ atom system. In Fig. 5.12a, we first show the late-time marginal probability of measuring a single-qubit subsystem A in the state $|0\rangle$, which is minimized at a detuning of 0.5 MHz, with a probability close to 0.5; this condition corresponds to the infinite effective temperature case. Similarly, we find the distance between the projected ensemble and the Haar-random ensemble (at the level of a 2nd order moment) also is

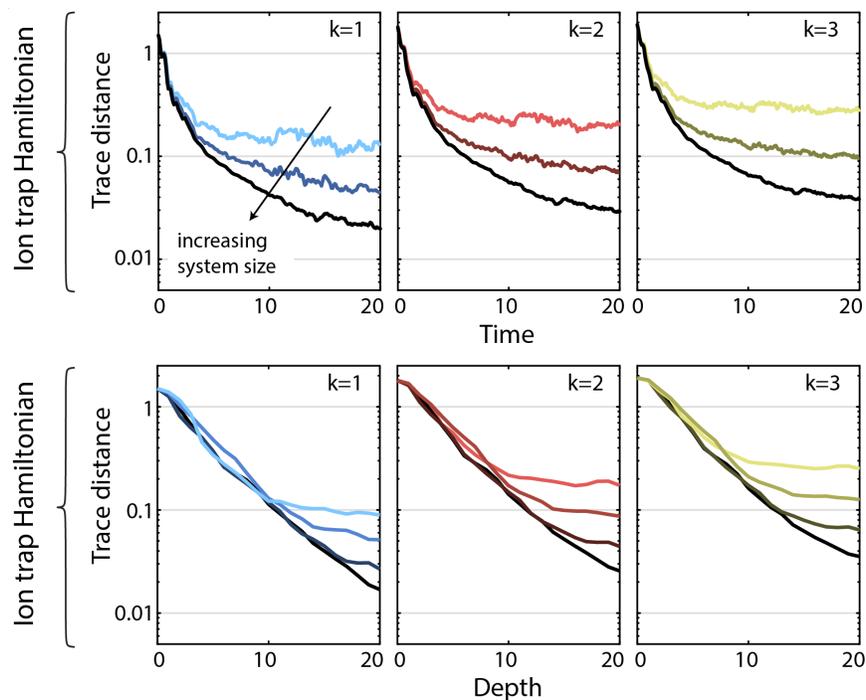


Figure 5.10: **Emergence of the projected ensemble in other quantum systems.** Numerically computed trace distances between the projected ensemble of a two-qubit subsystem and the corresponding k -design for RUC and trapped ion models. Results are shown for multiple different total system sizes: 10, 13, 16 for the trapped ion case, and 10, 12, 14, 16 for the RUC case, with darker colors corresponding to larger total system sizes.

minimized at this point (Fig. 5.12b). For three choices of detuning, indicated by vertical dashed lines in Fig. 5.12a and b, we plot the probability histograms for the single qubit subsystem, as in Fig. 5.5, and show the corresponding distribution of projected ensemble state vectors on the Bloch sphere (Fig. 5.12c).

We find that the numerical data in Fig. 5.12c agree well with an analytical form, $P_S(p) = (ap + b(1 - p))^{-3}$, where a and b are real constants fixed such that the distribution is normalized, $\int_0^1 dp P_S(p) = 1$, and the mean of the distribution corresponds to the marginal probability, $\int_0^1 dp P_S(p)p = p(z_A = 1)$ (red curves). This distribution, $P_S(p)$, is predicted from a so-called Scrooge ensemble of states, first proposed in Ref. [311] in the context of *quantum communication* (see Ref. [312] for a derivation of $P_S(p)$). It was later hypothesized that such an ensemble could appear in quantum chaotic dynamics [304].

The Scrooge ensemble consists of normalized states, $|\psi\rangle_\rho = \hat{\rho}^{1/2}|\psi\rangle / \sqrt{\langle\psi|\hat{\rho}|\psi\rangle}$, where $|\psi\rangle$ is drawn from a uniform, Haar-random ensemble with a weighing factor $D_S\langle\psi|\hat{\rho}|\psi\rangle$, where D_S is the dimension of the Hilbert space in which the Scrooge ensemble is defined. As its first-order moment, the Scrooge ensemble reproduces a desired target density operator,

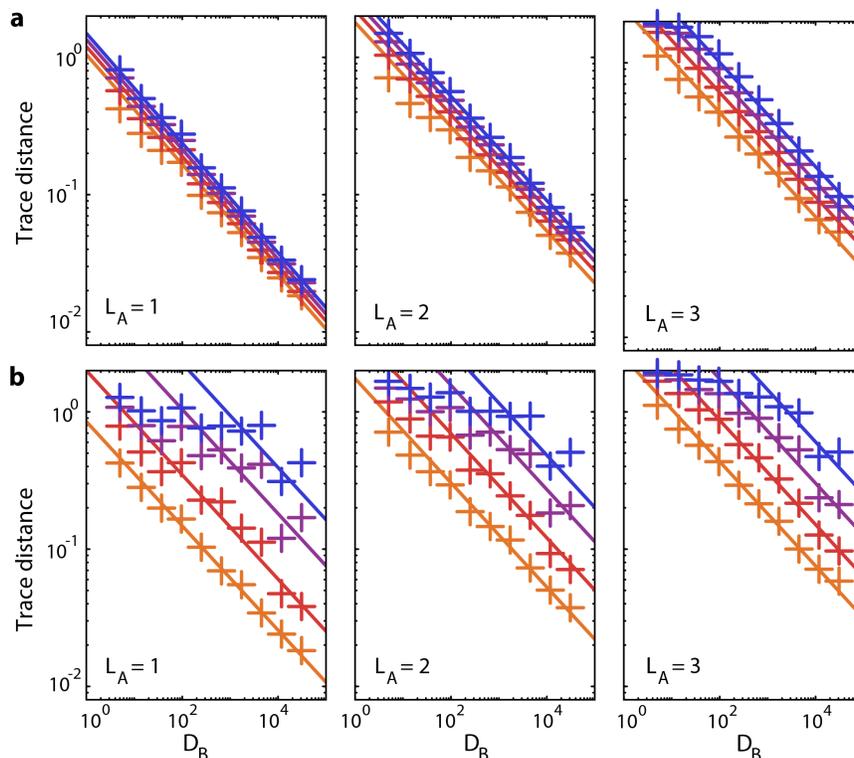


Figure 5.11: **Trace distance between the projected ensemble and quantum state k -designs:** **a, b**, System size scaling of late-time trace distances from three different subsystem sizes of $L_A = 1, 2$, and 3 . For fixed subsystem A size, we increase the Hilbert space dimension of outer subsystem B , D_B , and compute trace distances to the quantum state k -designs for four different orders, $k = 1$ (orange), 2 (red), 3 (purple), and 4 (blue). Sampling weighting factors of $p(z_B)$ and $p^k(z_B) / \sum_{z_B} p^k(z_B)$ are used in **a** and **b**, respectively, in the construction of a k th-order projected ensemble.

$\hat{\rho}$: $E_{\psi \in \text{Haar}}[D_S \langle \psi | \hat{\rho} | \psi \rangle | \psi \rangle_{\rho} \langle \psi |_{\rho}] = \hat{\rho}$; in our case, $\hat{\rho}$ refers to the reduced density operator of A . For the effective infinite effective temperature case ($\hat{\rho} = \hat{I}/D_A$), the Scrooge ensemble reduces to a homogeneous, Haar-random ensemble. Motivated by the the good agreement of probability histograms (Fig. 5.12c), we also plot the trace distance between the second moment of the Scrooge ensemble, $E_{\psi \in \text{Haar}}[D_S \langle \psi | \hat{\rho} | \psi \rangle (|\psi \rangle_{\rho} \langle \psi |_{\rho})^{\otimes 2}]$, and the second moment of the projected ensemble as a function of detuning (Fig. 5.12d). In contrast to the Haar ensemble, the distance to the Scrooge ensemble is now almost independent of detuning. Moreover, for a detuning that corresponds to finite effective temperature, we observe that the distance decreases monotonically with total system size (inset of Fig. 5.12d).

It is interesting to note that the projected ensemble can indeed be viewed through a quantum communication lens: projective measurements in B create the projected ensemble, which can be thought of as a message ensemble sent to an observer in A . Message ensembles are often characterized by the accessible information, the maximum information gain the

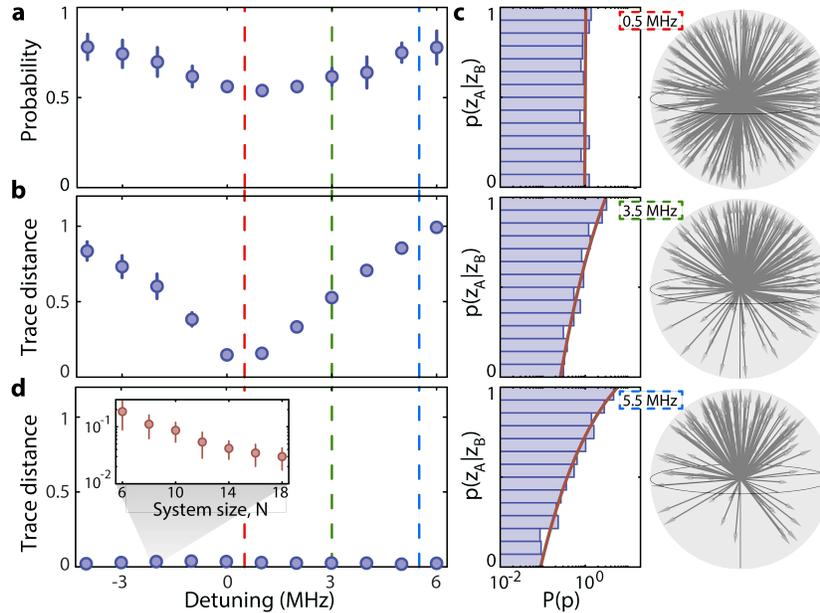


Figure 5.12: **Signatures of finite effective temperature projected ensembles.** **a**, Late-time, marginal probability values of $p(z_A = 0)$ as a function of global detuning for the Rydberg Hamiltonian with $N = 16$ atoms; error bars denote the standard deviation over temporal fluctuations. We find that at a detuning value of 0.5 MHz, the late-time probability value equilibrates around 0.5, signaling the infinite effective temperature condition where a uniformly random ensemble emerges. **b**, Trace distances between the projected ensemble and the Haar-random ensemble as a function of the global detuning; as in **a**, the distance is minimized at the infinite effective temperature condition of 0.5 MHz. **c**, Projected ensembles for a single-qubit subsystem and their conditional probability distributions at finite effective temperatures. Detuning values are chosen as 0.5, 3, and 5.5 MHz (indicated by dashed lines in **a** and **b**), respectively from top to bottom. Here, the subsystem A is taken from the middle atom at site 8. We find the numerical results are in good agreement with the analytic prediction from the so-called *Scrooge ensemble* (red lines, see text). **d**, Trace distances between the projected ensemble and the Scrooge ensemble as a function of the global detuning. The inset shows that the trace distance decreases monotonically as increasing total system size, N , which suggests that the single-qubit, finite-temperature projected ensembles converge to the Scrooge ensemble in the thermodynamic limit.

observer (by optimal measurement) can achieve about the message. Interestingly, Scrooge ensembles attain a lower bound of the accessible information if used as a message ensemble [311]. Our results in Fig. 5.12 suggest that Scrooge ensembles naturally appear in quantum many-body dynamics as projected ensembles (from a generic local basis measurement in B), hence minimizing the accessible information about the projected ensemble obtainable from measurements in A .

Again we recommend interested readers consult Ref. [71] for wealth of details on this subject, most of which we only briefly touch on here.

CHAPTER 6

Many-body fidelity benchmarking

While not a hard and fast rule, the conventional wisdom is that what makes quantum computers powerful is their ability to wield large entangled states. Entanglement¹ is a resource for quantum computation, simulation, and metrology, but in the modern NISQ [50] (noisy intermediate-scale quantum) era generating entanglement is far from a trivial task. By contrast, in these murky pre-error-correction days, experiments are beset by errors, limiting their fidelity of correctly performing a desired quantum evolution. However, the problem is even more fundamental than poor execution: except in special cases, we don't even know how to measure how badly we did!

In Chapter 3, we studied two particular instances of quantum evolution: preparation of two-atom Bell states, and preparation of long-range-ordered states of 26 atoms. In both cases, we were able to make quantitative statements as to our success probability, and in the two-atom case could even directly estimate the preparation fidelity. With such results in hand, it may seem we can simply skip forward, and apply similar techniques to the generic many-body case. Unfortunately, the very aspect of quantum science which makes it of such general interest – namely, its classical computational intractability – seemingly precludes us from efficiently understanding how well we can accomplish generic quantum tasks.

In this chapter, we will explore why we care about measuring quantum fidelities, how standard methods are unable to accomplish this task, and what our solution to this problem is. At the end, we will emerge with a general, efficient fidelity estimator, applicable for nearly any type of quantum evolution.

¹Throughout this thesis we will almost always use “entanglement” to imply an entropy-type quantity like von Neumann entanglement entropy or negativity. As we study in great detail, such extensive metrics are related to the complexity of representing quantum states. Thus when we say “highly entangled” we do NOT refer to classically representable states like Greenberger–Horne–Zeilinger (GHZ), cluster, or stabilizer states.

6.1 ESTIMATING FIDELITY

To get started, what do we even mean by fidelity in a many-body context? We note that much of the elementary discussion here can be found in much more detailed form in Ref. [15], Chapter 9; however, be warned that in Ref. [15] what they call “fidelity” is really the square root of the quantity we will be referring to, though our definition is the more standard modern nomenclature.

At its core, fidelity is a distance measure between quantum states, and quantifies their overlap. While fidelity can be defined between two mixed states, in this thesis we will always be interested in the fidelity overlap between a mixed state, $\hat{\rho}$, (generally generated by the experiment) and a pure state, $|\psi\rangle$, which allows us to simply write the fidelity, F , as

$$F = \langle \psi | \hat{\rho} | \psi \rangle. \quad (6.1)$$

Intuitively, you can think of the fidelity as a roughly a success likelihood, i.e. it is the probability that when we execute a quantum operation we correctly generate the pure state we desire. Thus it is a global quantity measuring the “goodness” of our operation.

Besides just being a success metric, however, the fidelity also encodes information about how well we can estimate any observable. With some math expounded upon in more detail in Ref. [15], we can write

$$\max_{\{\hat{E}_j\}} \frac{1}{2} \sum_j |p_j - q_j| \leq \sqrt{1 - F}, \quad (6.2)$$

where $\{\hat{E}_j\}$ is a POVM, $p_j = \text{Tr}(|\psi\rangle\langle\psi|\hat{E}_j)$, and $q_j = \text{Tr}(\hat{\rho}\hat{E}_j)$. The l.h.s. of Eq. (6.2) is the cumulative absolute difference between observable probability distribution for a given set of measurements on the quantum state, maximized over all possible choices of POVM observables. Thus, more plainly, Eq. 6.2 states that knowledge of F allows us to upper bound the *maximum possible variation* between observables of the experimental mixed state and the target pure state.

Unfortunately, however, measuring the fidelity is a tricky business for all but the smallest quantum systems for two reasons. One, it is in general hard to predict what $|\psi\rangle$ is in generic situations where classical simulations become intractable, and two, it is exceedingly difficult to reconstruct $\hat{\rho}$ from experimental measurements. The latter task is known as quantum state tomography [313–315], and in its most naïve implementation it requires an exponential number of measurements. This is because it requires measurement in an exponential number of bases to have access to all 4^N components of $\hat{\rho}$, and requires an exponential number of

measurements in each basis to accurately reconstruct the exponentially small terms. In other words, it is practically impossible for more than a few qubits.

6.1.1 The linear cross-entropy benchmark

Clearly if our goal is to measure the fidelity, quantum state tomography is a non-starter, and we will need to take a different approach. Luckily significant theoretical work [275, 276] was devoted exactly to this question, and a solution was uncovered: employ randomness in quantum evolution to estimate the fidelity with high accuracy with few measurements. We highlight Ref. [297] for a very pedagogical and beginner-friendly introduction to many topics we shall discuss.

The idea essentially is the following. To start, begin with your system in a simple to prepare initial state. Then apply sufficiently randomizing, or scrambling, evolution such that the quantum state ergodically explores its Hilbert space. After some finite time, measure the state of the system in a fixed basis – call it the Z -basis for simplicity – to resolve the probability distribution, $q(z)$. At the same time, simulate the same exact evolution on a classical computer with no errors, yielding a probability distribution $p(z)$. Now, calculate the following quantity

$$F_{XEB} = (D + 1) \left(\sum_z q(z)p(z) \right) - 1, \quad (6.3)$$

where $D = 2^N$ is the Hilbert space dimension. If $q(z) = p(z)$, i.e. if the distributions are perfectly correlated, then the sum becomes the second moment of $p(z)$. While this is generically unhelpful, the key thing to note is that we chose our dynamics to be highly scrambling, meaning the output of the experiment should look as if it was essentially drawn from a random state ensemble – the same random state ensembles we have discussed *ad nauseum* in Chapter 5. Under this assumption, we know that the second moment should be equal to $\sum_z p(z)^2 = 2/(D + 1)$, such that the r.h.s of Eq. (6.3) simplifies to 1.

Now, consider the opposite limit, where the experiment experiences so many errors as to reduce it to a maximally mixed state, $\hat{\rho} = \hat{I}/D$, such that $q(z) = 1/D$ for all z . In this case, the two probability distributions are completely uncorrelated, and Eq. (6.3) simplifies to $1/D$. Thus at least for these two limits of perfect correlation or complete uncorrelation, the quantity F_{XEB} appears to mimicking what we expect as the fidelity between the experiment and the ideal.

We can go a step further, and make this connection more concrete. Let us assume the mixed

state can be written as

$$\hat{\rho} = F|\psi\rangle\langle\psi| + (1 - F)\hat{\sigma}, \quad (6.4)$$

where $\hat{\sigma}$ absorbs all the error in the system and we define $\langle\psi|\hat{\sigma}|\psi\rangle = 0$ such that F is the fidelity¹. We further define $s(z) = \langle z|\hat{\sigma}|z\rangle$, in which case we write the following proof:

Proof.

$$\begin{aligned} F_{XEB} &= (D + 1)\left(\sum_z q(z)p(z)\right) - 1 \\ &= (D + 1)\left(\sum_z (Fp(z) + (1 - F)s(z))p(z)\right) - 1 \\ &= (D + 1)\left(F\sum_z p(z)^2 + (1 - F)\sum_z s(z)p(z)\right) - 1 \\ &= (D + 1)\left(\frac{2F}{D + 1} + \frac{(1 - F)}{D}\right) - 1 \\ &= F + \frac{1 - F}{D} \\ &\approx F, \end{aligned} \quad (6.5)$$

where in the last line we have taken $D \rightarrow \infty$, and where crucially we have assumed that $s(z)$ is uncorrelated with $p(z)$.

Thus we have essentially shown that – assuming the target pure state is sufficiently random – F_{XEB} is a good estimator for the many-body fidelity.

6.1.2 Caveats of cross-entropy benchmarking

There are some obvious subtleties with Eq. (6.5), which will motivate our development of more optimal fidelity estimators in the next sections.

Anti-correlation

For one, here we have only considered the case where $s(z)$ and $p(z)$ are well-correlated or uncorrelated. However, it is of course equally valid that the two distributions could be *anti-correlated*, for instance when they have no mutual support (e.g. let $s(z) = \delta(z - z_1)$ and $p(z) = \delta(z - z_2)$ where δ is the Dirac-delta function, and $z_1 \neq z_2$). In this case $F_{XEB} < 0$,

¹Oftentimes, we take $\hat{\sigma} = \hat{I}/D$, i.e. it is a maximally mixed state. In this case, however, F is not quite the fidelity as we typically define it at the level of $1/D$, and one must properly account for the overlap between $|\psi\rangle$ and $\hat{\sigma}$.

which is clearly unphysical behavior for estimating the fidelity which is strictly non-negative. From extensive numerical evidence, this case is atypical for normal experimental noise sources. However, it shall rear its head dramatically in Chapter 8 when we work with approximate classical simulations of the quantum dynamics, and we shall employ innovative methods to rein it in.

Invisibility of errors

Furthermore, the assumption that $s(z)$ and $p(z)$ are uncorrelated is partially violated in many other cases. In the most pathological and unrealistic scenario, if an error occurs which commutes both with the Hamiltonian and the projectors into the measurement basis, it will be wholly undetectable.

Much more realistically, $s(z)$ and $p(z)$ can remain correlated at very short times after an error occurs. For instance, if a phase flip error occurs, it is not immediately visible in the fixed measurement basis, and so $s(z)$ and $p(z)$ could still be highly correlated even though the global many-body fidelity has diminished. However, for sufficiently scrambling dynamics, the error will again become visible after a short *delay time*.

To understand why, it is easiest to work in the Heisenberg picture. In this case, we can consider what happens to the operator \hat{Z}_i (a phase flip occurring on site i) as a function of time. Under action of the unitary operator \hat{U} , we can write

$$\hat{U}\hat{Z}_i\hat{U}^\dagger = \sum_{\mu_1\mu_2\dots\mu_N} A_{\mu_1\mu_2\dots\mu_N} \hat{\mu}_1\hat{\mu}_2\dots\hat{\mu}_N, \quad (6.6)$$

where the coefficient A on the r.h.s. is the quantum amplitude of a specific Pauli string operator, where the $\hat{\mu}_i$ are Pauli operators drawn from $\{\hat{X}, \hat{Y}, \hat{Z}, \hat{I}\}$.

Essentially, Eq. (6.6) says that an initially localized error can transform into a highly non-local string operator under unitary evolution. The fundamental assumption of “sufficient scrambling” is that the coefficients A are all roughly equal, such that errors in any basis will show some signature in the measurement basis, albeit potentially after a short delay.

We can model this effect by assuming $s(z) = p(z) \exp(-t/\tau) + (1 - \exp(-t/\tau))/D$, where τ is the typical delay time. Inserting this expression into Eq. (6.5), we find

$$\begin{aligned} F_{XEB} &= (D+1) \left(\frac{2F}{D+1} + (1-F) \sum_z s(z)p(z) \right) - 1 \\ &= (D+1) \left(\frac{2F}{D+1} + (1-F) \sum_z \left(p(z) \exp(-t/\tau) + \frac{1 - \exp(-t/\tau)}{D} \right) p(z) \right) - 1 \\ &\approx F(1 - \exp(-t/\tau)) + \exp(-t/\tau), \end{aligned} \quad (6.7)$$

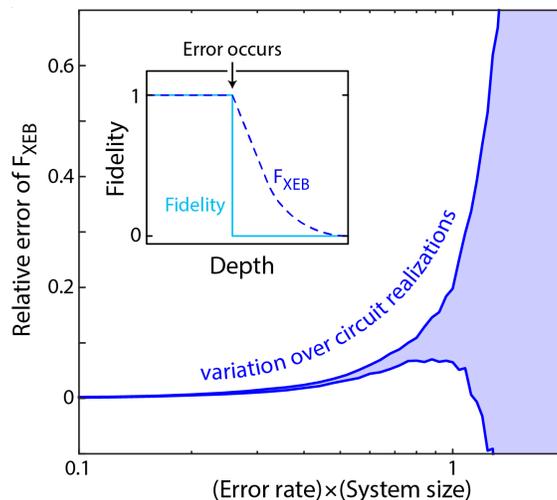


Figure 6.1: **Breakdown of fidelity estimation.** Inset: When an error occurs, the fidelity estimator (for instance F_{XEB}) does not instantly follow the decrease in fidelity, but only does so after a finite delay time. Main plot: fidelity estimation fails beyond a certain error rate, as errors accumulate faster than they are recognized due to the delay time. This breakdown occurs roughly when the total error probability per layer approaches unity. Here this is visualized for $N = 10$ RUC dynamics with random single-qubit Pauli errors, where the breakdown occurs around an independent single qubit error rate of 0.1.

where in the last line we have again taken the limit $D \rightarrow \infty$. Thus, F_{XEB} only is an accurate estimator of F after a delay time of τ has elapsed (Fig. 6.1 inset).

We note that this delay time has more dramatic ramifications than just making fidelity estimation slightly inaccurate. In particular, consider the case where the error rate grows large enough that on average more than one error occurs per circuit depth. In this case, when an error occurs the fidelity estimator will start dropping (with a characteristic timescale of the delay time) to match the true fidelity, but before it has equilibrated another error will occur, affecting the fidelity. This will lead to a staggered process where the fidelity estimator becomes entirely inaccurate to the true fidelity.

We visualize this process in Fig. 6.1, where we see for $N = 10$ RUC dynamics with single-qubit Pauli errors the relative mis-estimation of the fidelity estimator remains small and relatively flat until the error rate is 0.1, at which point the fidelity estimator becomes completely erroneous. This breakdown point occurs roughly where more than one error occurs per circuit layer, and is thus sensitive to both the number of qubits, and the error rate. We note that this observation is intimately linked to other studies of phase transitions in cross-entropy benchmarking [316–318], but recasts it in a more intuitive light.

Necessity of global scrambling

Finally, F_{XEB} is only a good description if the global state is as random as a 2-design, i.e. if the second moment $p(z) = 2/(D+1)$. There are various reasons why this may not be the case. For instance, at the earliest time, let $q(z) = p(z) = \delta(z - z_0)$, where $|z_0\rangle$ is the initial fiduciary product state. Then we would have $F_{XEB} = (D+1) \sum_z \delta(z - z_0)^2 - 1 = D+2$, which is exponentially larger than 1, and would slowly diminish to match the actual fidelity as scrambling was ongoing (see for instance Ref. [316]). Still, the F_{XEB} is not a fidelity estimator until the global second moment is that of a Haar random state, which is limited by the Lieb-Robinson velocity [319, 320].

Furthermore, this condition is impossible if the system is constrained by conservation laws or symmetries such that a global 2-design is never achieved. As an example, imagine a finite temperature scenario such that the $|1\rangle$ state is far preferred over the $|0\rangle$ state. In this case, the second moment cannot reach the needed value. However, as we have explored extensively in Chapter 5, there are still ways to identify random state ensemble behavior even in cases where global scrambling has not occurred due to Lieb-Robinson bounds, symmetries, or conservation laws.

In the next section, we will be employing these innovative types of random state ensembles to solve, or at least partially address, many of the insufficiencies of F_{XEB} . We will proceed in two steps, and introduce two new fidelity estimators, F_c and F_d . F_c is designed to make the standard linear cross-entropy perform better at early times via the use of *local scrambling* from the projected ensemble (see Chapter 5 for details thereof). F_d then will build off of F_c , maintaining its same benefits while opening them up to many new types of quantum systems.

6.2 BUILDING A BETTER CROSS-ENTROPY

As just discussed, the linear cross-entropy benchmark suffers from several problems, such as failure at early times and lack of applicability to systems with conservation laws or symmetries. In this section, we will tackle these problems, and introduce a new, more general cross-entropy which estimates the fidelity both at short and early times, and is applicable to a wide range of quantum dynamics. We note that especially for the latter point, the definitive reference is Ref. [30], which should be consulted for extensive details.

A key question is if the formation of approximate k -designs in the projected ensemble enables associated applications like device benchmarking with only global, time-independent control. Decoherence changes observables of the projected ensemble (Figs. 5.6); can this

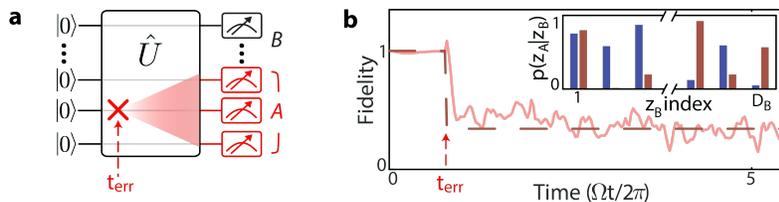


Figure 6.2: **Detecting errors during quantum evolution.** **a**, Schematic of noisy time evolution with an error occurring at time t_{err} . The influence of the local error propagates outward, affecting the measurement outcomes non-locally at a later time. **b**, Errors during evolution can be detected by correlating the measurement outcomes with an error-free, ideal evolution case. We numerically tested this by applying a local, instantaneous phase error to the middle qubit of an $N = 16$ atom Rydberg simulator at time $\Omega t_{\text{err}}/2\pi \approx 1$. The proposed fidelity estimator, F_c (solid line), accurately approximates the many-body overlap (dashed line) between states produced with and without errors, after a slightly delayed time. Inset: Conditional probability distributions in A before (blue) and after (red) the error, showing decorrelation.

quantitatively signal the onset of decoherence in a quantum device?

We affirmatively answer this question by benchmarking the evolution of our experimental system under a time-independent Hamiltonian. As a toy example, we consider the case of a single error occurring at time t_{err} during unitary evolution. The effect of this error then propagates outward [320], generically transforming the evolution output state and affecting measurement outcomes in subsystem A (Fig. 6.2). Using the fact that the projected ensemble forms an approximate 2-design (see Chapter 5), we devise a fidelity estimator F_c to quantify the effect of this error as a rescaled cross-correlation between measurement probabilities in the experimental and ideal conditions:

$$F_c = 2 \frac{\sum_z p(z)q(z)}{\sum_z p^2(z)} - 1, \quad (6.8)$$

where $q(z)$ and $p(z)$ are the experimental and theoretical probabilities of observing a global bitstring z , respectively. We numerically confirm that shortly after we apply an instantaneous phase rotation error on one qubit, our estimator approximates the many-body state overlap, $F_c \approx F = \langle \psi | \hat{\rho} | \psi \rangle$, between the ideal state, $|\psi\rangle$, and the erroneous state, $\hat{\rho}$ (Fig. 6.2b).

To evaluate F_c experimentally, we formulate an empirical, unbiased estimator:

$$F_c \approx 2 \frac{\frac{1}{M} \sum_{i=1}^M p(z_{\text{exp}}^{(i)})}{\sum_z p^2(z)} - 1, \quad (6.9)$$

where M is the number of measurements and $z_{\text{exp}}^{(i)}$ is the experimentally measured bitstring at the i th repetition. While this reformulation still requires calculation of a reference

theory comparison, we note that the required number of experimental samples to accurately approximate F_c scales favorably with system size N . Concretely, the standard deviation of F_c is estimated to be $\sigma(F_c) \approx 1.04^N/\sqrt{M}$ (Fig. 6.6), yielding an improved sample complexity in comparison to other existing methods [321]. This implies that we do not need to fully reconstruct the experimental probability distribution for fidelity estimation of large quantum systems.

6.2.1 Experimental results

We test our benchmarking protocol for errors occurring continuously with a Rydberg quantum simulator of up to $N = 25$ atoms. We estimate the fidelity of our experimental device, $F_{c,\text{exp}}$, by correlating measured bitstrings to results from error-free simulation as a function of evolution time. In addition, we use an *ab initio* error model with no free parameters that mimics the experimental output, from which we extract both the fidelity estimator, $F_{c,\text{model}}$, and the model fidelity, $F_{\text{model}} = \langle \psi(t) | \hat{\rho}_{\text{model}}(t) | \psi(t) \rangle$ (Fig. 6.3a).

In Fig. 6.3b, we compare F_{model} , $F_{c,\text{exp}}$, and $F_{c,\text{model}}$ for a system of ten atoms. We observe $F_{c,\text{model}} \approx F_{\text{model}}$, validating the efficacy of the estimator under realistic error sources. Additionally, we find $F_{c,\text{exp}} \approx F_{c,\text{model}}$, and that full bitstring probability distributions show good agreement between the error model and the experiment, indicating that our *ab initio* error model is a good description of the experiment.

We further apply this method to estimate the fidelity for generating states with a maximum half-chain entanglement entropy in larger systems. We first use error-free simulation to calculate the half-chain entanglement entropy growth as a function of system size, finding that the entanglement saturates at a time, t_{ent} , linear in system size (Fig. 6.3c). We evaluate the fidelity estimator F_c for N ranging from 10 to 25, each at their respective t_{ent} , again finding good agreement between experiment and our *ab initio* error model in the range for which our error model is readily calculable (Fig. 6.3d). We note an estimated fidelity of 0.49(2) for generating a state with maximum half-chain entanglement entropy for $N = 25$.

We numerically show F_c also applies for erroneous evolution using other quantum devices (Fig. 6.4). For both random unitary circuits and Hamiltonian evolution, F_c accurately estimates the fidelity at much shorter evolution times than do existing methods such as linear cross-entropy benchmarking [162, 276], which can be explained through the early time-formation of the projected ensemble. Specifically, the Hamiltonian system we simulate

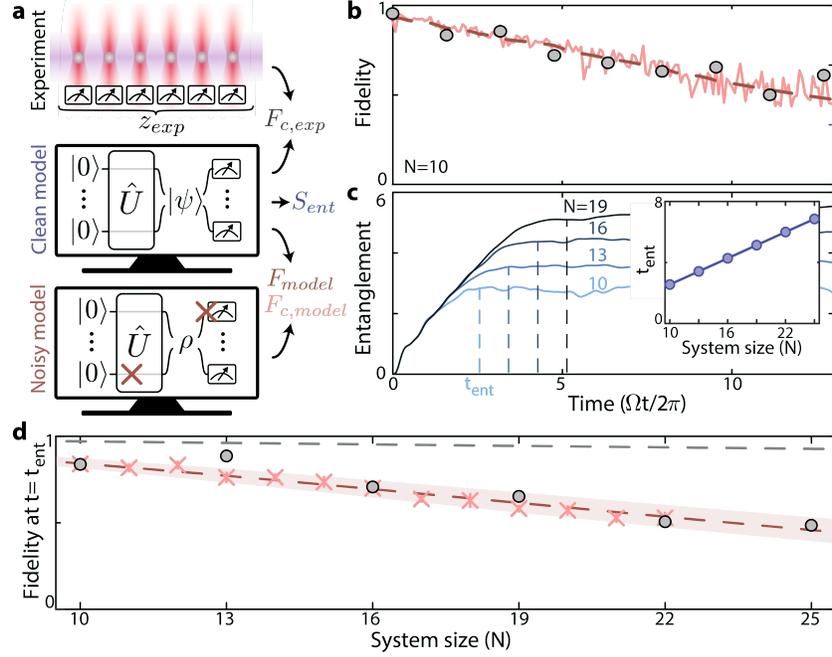


Figure 6.3: **Fidelity estimation of an analog Rydberg quantum simulator.** **a**, To estimate experimental fidelity, we repeatedly perform Hamiltonian evolution, each time performing a projective measurement to accrue an ensemble of measured bitstrings z_{exp} . We then correlate the measured bitstrings with an error-free simulation of the dynamics in order to calculate the fidelity estimator, $F_{c,exp}$. To validate our fidelity estimation method, we compare the error-free simulation against results from an *ab initio* error model, to calculate the model fidelity F_{model} and accompanying estimator $F_{c,model}$. **b**, Experimental benchmarking of a Rydberg quantum simulator for $N=10$ atoms with blockaded Hilbert space dimension $D=144$. Shown are $F_{c,exp}$ (grey markers), the fidelity F_{model} (dashed red line), and $F_{c,model}$ (solid pink line). **c**, The half-chain entanglement entropy (calculated from the error-free simulation) increases before saturating at a time, t_{ent} , which grows linearly with system size (inset). **d**, Fidelity estimated at t_{ent} , showing estimator F_c from experiment (grey markers) up to $N=25$, and from error model (pink crosses) up to $N=22$. Additionally, we show a fit to the model fidelity, given as $F_0^N \exp[-\gamma(N)t_{ent}(N)]$ (red dashed line), where F_0 is the single-atom preparation fidelity and $\gamma(N)$ is the many-body fidelity decay rate of our Rydberg simulator. The fidelity estimation uses only $<10^4$ experimentally sampled bitstrings per data point.

is a trapped ion Hamiltonian:

$$\hat{H}_{ion}/\hbar = 2\pi J \left[\sum_{i=1}^N 0.4\hat{S}_i^x + 0.45\hat{S}_i^y + \sum_{i=1}^{N-1} \sum_{j=i+1}^N \frac{\hat{S}_i^x \hat{S}_j^x}{|i-j|} \right], \quad (6.10)$$

where $\hat{S}_i^{x,y,z}$ are the spin-1/2 operators at site i . The Hamiltonian coefficients are adapted from Ref. [322], including the $1/r$ long-range interactions. In addition, coefficients are globally rescaled by J such that the early-time growth rate of half-chain entanglement

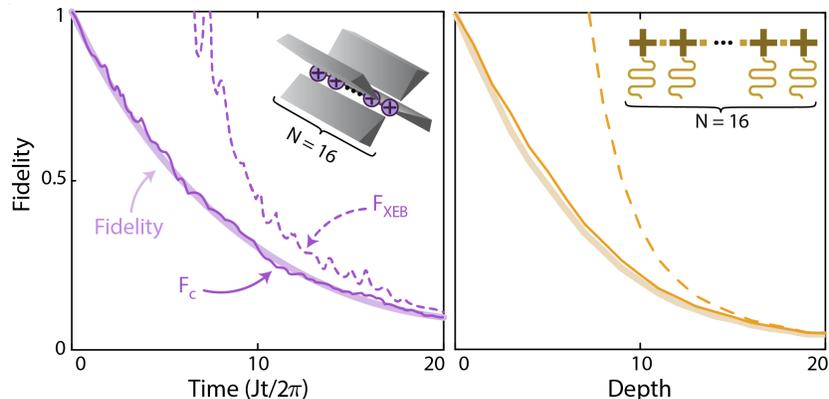


Figure 6.4: **Benchmarking for other quantum systems.** **a**, Fidelity estimation for the case of a trapped ion quantum simulator governed by chaotic Hamiltonian evolution (left) and a quantum computer implementing a random unitary circuit (RUC) (right). In both cases, we plot the many-body fidelity (dashed line), as well as our fidelity estimator, F_c (solid line) and the more conventional linear cross-entropy-benchmark, F_{XEB} [162] (dotted line). We find that F_c approximates the fidelity at much earlier times than F_{XEB} .

entropy is matched to that of RUC dynamics. The many-body quantum state at time t is obtained by solving for $|\psi(t)\rangle = e^{-i\hat{H}_{\text{ion}}t/\hbar}|0\rangle^{\otimes N}$, and noisy evolution is represented by local bit-flip and phase-flip errors. Random unitary circuit details are summarized in Appendix B.

6.2.2 Data analysis

Our state readout is described in detail in Ref. [34]; it features single-site detection which discriminates atoms in the clock state, $|0\rangle$, versus the Rydberg state, $|1\rangle$, through a combination of fluorescence imaging and Rydberg auto-ionization. We take a total of three images: 1) after the array is initially loaded to perform rearrangement, 2) after the rearrangement is completed to verify the initial state is correct, 3) after the sequence has finished. We post-select for image triplets where the proper rearrangement pattern is visible in image (2), and calculate the survival of each atom by comparing site occupations in image (2) to image (3).

This array of survival signals is then converted into the qubit basis. For instance, in typical experiments where atoms are rearranged into defect-free arrays of ten atoms, we calculate the binary survivals for each atom, and then make the mapping ‘atom survived’ $\rightarrow |0\rangle$ and ‘atom did not survive’ $\rightarrow |1\rangle$, yielding a bitstring of the qubit states. After taking many shots we accrue an ensemble of such bitstrings, $\{z\}$. For randomness measurements, a total of ≈ 120000 shots are used (≈ 3000 shots per time point). For benchmarking measurements a total of ≈ 40000 shots are used for generating the time-trace at $N = 10$ in Fig. 6.3b (≈ 3700

shots per time point). Approximately ≈ 44000 total shots are used for the N -scaling plot in Fig. 6.3d, where the number of shots for a given system size is approximately given by $M \approx 3000 + 250N$.

In Fig. 6.3, error bars on experimental quantities are calculated via extrapolation from subsampling of the total number of experimentally measured bitstrings to estimate the sample complexity at a given N (Fig. 6.6). The error bars on $F_{c,\text{model}}$ from the *ab initio* error model stem from typicality errors associated with the temporal fluctuation of our estimator.

6.2.3 Calculating the entanglement saturation time

As can be seen in Fig. 6.3c, entanglement growth in our Rydberg quantum simulator is generally characterized by two distinct regions: a size-independent linear increase, followed by saturation at an N -dependent value. In order to systematically capture this behavior and predict the entanglement saturation time for arbitrary N , we apply the following protocol. We first calculate the entanglement growth for system sizes ranging from $N = 10$ to 22. We then fit all profiles with a functional form of

$$S_{\text{ent}}(N, t) = \begin{cases} m_1 t & t \leq t_c(N) \\ m_1 t_c(N) + m_2(N)(t - t_c(N)) & t > t_c(N) \end{cases}$$

with free parameters $m_1, m_2(N)$, and $t_c(N)$, but with the explicit constraint that m_1 must be the same for all system sizes. From this we find $t_c(N)$, and we further define $t_{\text{ent}}(N) = Ct_c(N)$, where in the Rydberg case we set $C = 1.35$ in order to make sure the time we study is firmly in the saturated entanglement regime (as can be verified visually in Fig. 6.3c). The secondary slope m_2 is used because even past t_{ent} there is still some slight growth to the entanglement entropy, which becomes more noticeable for larger N ; we explore this in slightly more detail in Section 9.2.1. The entanglement saturation time is then fit as a linear function of system size, yielding $t_{\text{ent}}(N) = \alpha_0 + \alpha_1 N$; for our particular Hamiltonian parameters we find $\alpha_0 = -0.0580(2)$, $\alpha_1 = 0.05404(1)$, both in units of μs , or $\alpha_0 \approx -0.3$ and $\alpha_1 \approx 0.3$ in units of cycles.

6.2.4 Derivation of the fidelity estimator

Here we prove the fidelity estimator F_c is a good estimator of the fidelity, in the sense that it captures the effect of single errors supported in local subsystems; by averaging over many such local subsystems, the global fidelity can be captured even at early times.

Proof. Our fidelity estimator F_c (Eq. 6.8) can be understood by expressing the global bitstring probabilities for ideal and noisy evolution, $p(z)$ and $q(z)$ respectively, in terms of conditional and marginal probabilities as

$$p(z) = p(z_A|z_B)p(z_B) \quad (6.11)$$

$$q(z) = q(z_A|z_B)q(z_B), \quad (6.12)$$

for complementary subsystems A and B . We consider the simplest case of a single local error \hat{V} occurring at time t_{err} during time-evolution, and assume that the time-evolved error operator, $\hat{V}(\tau) = \hat{U}(\tau)\hat{V}\hat{U}(\tau)^\dagger$, is supported within subsystem A . Here $\tau = t - t_{\text{err}}$ is the time past the occurrence of the error and $\hat{U}(\tau)$ is the time-evolution operator from t_{err} to t . This implies that the measurement outcome in B is not affected by the error, giving $q(z) = q(z_A|z_B)p(z_B)$ because $q(z_B) = p(z_B)$. Under these conditions, we can rewrite F_c as

$$F_c = 2 \frac{\sum_z p(z)q(z)}{\sum_z p(z)^2} - 1 \quad (6.13)$$

$$= 2 \frac{\sum_{z_B} p^2(z_B) \sum_{z_A} q(z_A|z_B)p(z_A|z_B)}{\sum_{z_B} p^2(z_B) \sum_{z_A} p^2(z_A|z_B)} - 1 \quad (6.14)$$

$$\approx \frac{\sum_{z_B} p^2(z_B)(F_{\text{XEB}}(z_B) + 1)}{\sum_{z_B} p^2(z_B)} - 1 \quad (6.15)$$

$$= \sum_{z_B} r(z_B) F_{\text{XEB}}(z_B) \quad (6.16)$$

where $r(z_B) = \frac{p^2(z_B)}{\sum_{z_B} p^2(z_B)}$ and

$$F_{\text{XEB}}(z_B) = (D_A + 1) \sum_{z_A} q(z_A|z_B)p(z_A|z_B) - 1 \quad (6.17)$$

is the z_B -dependent, linear cross-entropy benchmark [162] in subsystem A , and D_A is the Hilbert space dimension of A . From Eq. (6.14) to Eq. (6.15), we used the second-order moment of the projected ensemble in an error-free case

$$\frac{1}{D_A} \sum_{z_A} p^2(z_A|z_B) \approx \frac{2!}{D_A(D_A + 1)} \quad (6.18)$$

based on our experimental and numerical observations of emergent local randomness during chaotic quantum dynamics (Chapter 5.2).

The validity of the relation $F_c \approx F$ can be analytically understood based on the assumption that the projected ensemble of $|\psi_A(z_B)\rangle$ approximately forms a quantum

state 2-design. To see this explicitly, we consider

$$\begin{aligned}
& \sum_{z_B} r(z_B) p(z_A|z_B) p(z_A|z_B) \\
&= \sum_{z_B} r(z_B) \langle \psi_A(z_B) | \hat{P}_{z_A} | \psi_A(z_B) \rangle \langle \psi_A(z_B) | \hat{V}^\dagger(\tau) \hat{P}_{z_A} \hat{V}(\tau) | \psi_A(z_B) \rangle \\
&= \text{tr} \left\{ \left(\hat{P}_{z_A} \otimes \hat{V}^\dagger(\tau) \hat{P}_{z_A} \hat{V}(\tau) \right) \cdot \sum_{z_B} r(z_B) (|\psi_A(z_B)\rangle \langle \psi_A(z_B)|)^{\otimes 2} \right\} \\
&\approx \frac{\text{tr} \left\{ \left(\hat{P}_{z_A} \otimes \hat{V}^\dagger(\tau) \hat{P}_{z_A} \hat{V}(\tau) \right) \cdot (\hat{\mathbb{1}} + \hat{\mathcal{S}}_A) \right\}}{D_A(D_A + 1)} \\
&= \frac{1 + \left| \langle z_A | \hat{V}(\tau) | z_A \rangle \right|^2}{D_A(D_A + 1)} \tag{6.19}
\end{aligned}$$

where $\hat{P}_{z_A} = |z_A\rangle \langle z_A|$ is the z -basis projector onto a specific bitstring z_A in A , $r(z_B)$ is the probability weighting factor, $\hat{\mathbb{1}}$ is the identity operator, and $\hat{\mathcal{S}}_A$ is the swap operator acting on subsystem A for the duplicated Hilbert space $\mathcal{H}_A^{\otimes 2}$. In order to obtain the fourth line, we used

$$\sum_{z_B} r(z_B) (|\psi_A(z_B)\rangle \langle \psi_A(z_B)|)^{\otimes 2} \approx \frac{\hat{\mathbb{1}} + \hat{\mathcal{S}}_A}{D_A(D_A + 1)}, \tag{6.20}$$

where the right-hand side is due to the projected ensemble forming an approximate quantum state 2-design. We note that the weighting factors $q(z_B)$ are different than those used for the majority of the thesis; however, we numerically find that approximate 2-designs form regardless of which weighting factor is used.

Inserting Eq. (6.19) into Eq. (6.16), we obtain

$$\begin{aligned}
F_c &\approx \frac{1}{D_A} \sum_{z_A} \left| \langle z_A | \hat{V}(\tau) | z_A \rangle \right|^2 \\
&= \frac{1}{D} \sum_z \left| \langle z | \hat{V}(\tau) | z \rangle \right|^2, \tag{6.21}
\end{aligned}$$

where the equality on the second line holds because one can always multiply the identity $\frac{1}{D_B} \sum_{z_B} \langle z_B | z_B \rangle^2 = 1$ with the Hilbert space dimension of the complement $D_B = D/D_A$ with D being the Hilbert space dimension of the entire system.

The relation in Eq. (6.21) explains how F_c estimates the many-body fidelity with a good accuracy. The right-hand side of Eq. (6.21) describes the return probability of $\hat{V}(\tau)$ (also

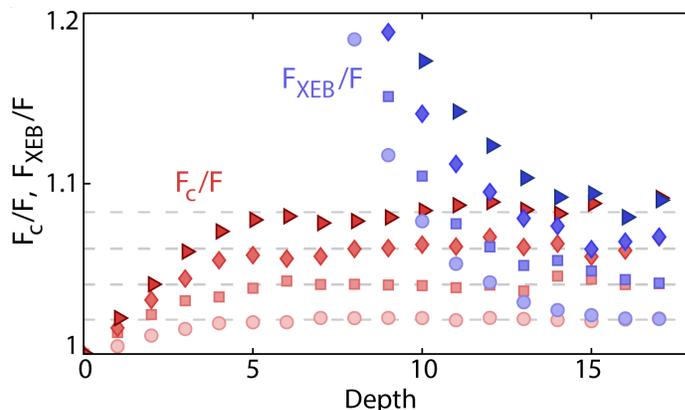


Figure 6.5: **Systematic delay time in fidelity estimation. a**, Systematic error in the fidelity estimator F_c (red) and the linear cross-entropy benchmark F_{XEB} (blue) relative to the true fidelity F for a ten-qubit random quantum circuit. Noisy quantum dynamics are simulated with random, local bit-flip and phase-flip errors at four different single-qubit error rate per circuit depth, $\gamma_{err} = 0.0025$ (circle), 0.005 (square), 0.0075 (diamond), and 0.01 (triangle). The dashed grey lines are the phenomenological error scaling of $\exp(N\gamma_{err}\tau_d)$ with an γ_{err} -independent delay time of $\tau_d = 0.8$ and $N = 10$. The saturation of both F_c/F and F_{XEB}/F at large circuit depth implies that at late times there is a delayed response in fidelity estimation performed with a fixed measurement basis, i.e. $F = \exp(-N\gamma_{err}t)$, but $F_c = \exp(-N\gamma_{err}(t - \tau_d))$.

known as Loschmidt echo) averaged over all possible initial states in the fixed measurement basis. Under chaotic time evolution, the propagated error operator $\hat{V}(\tau)$ becomes scrambled, and it is exponentially unlikely in the size of A that a computational state remains unchanged.

Therefore, non-vanishing contributions to F_c arise only when the error operator is partly proportional to the identity, e.g. $\hat{V}(\tau) = c_0\hat{\mathbb{1}} + \sum_s c_s(\tau)\hat{\sigma}_s$ with $c_0 \neq 0$, where s enumerates over all possible Pauli strings supported in A [320]. In such a case, $F_c \approx |c_0|^2$ approximates the probability that \hat{V} did not affect the many-body wavefunction, hence $F_c \approx F$. This statement becomes exact if the local qubit on which the error occurs is maximally entangled with the rest of the system at the time of the error. Our analysis can be straightforwardly generalized to more than one error, either located nearby or distant, as long as their joint support A leads to a random ensemble approximately close to the state 2-design.

Finally, we comment on the conditions in which F_c may significantly deviate from F . If \hat{V} is diagonal in the measurement basis, e.g. dephasing error along the z axis, and if the error occurs shortly before the bitstring measurements, the return probability in Eq. (6.21) will be close to unity despite that the many-body fidelity may be decreased significantly. Our method can fail in this special case. However, if F_c is evaluated after some delay time from the error (Fig. 6.5), then $\hat{V}(\tau)$ becomes scrambled in the operator basis, and F can be

approximately estimated (Fig. 6.2). In other words, even in the case of off-diagonal errors, our formula becomes valid after a finite delay time.

6.2.5 Fidelity estimation in the Rydberg model

Bitstring probability distributions in our system are characterized in a constrained Hilbert space \mathcal{H}_c where the simultaneous excitation of two neighboring atoms to the Rydberg state is not allowed. In order to benchmark our quantum device in the subspace \mathcal{H}_c , we modify the fidelity estimator F_c as

$$F_c(t) = B(t)B_0(t) \left(2 \frac{\sum_{z \in \mathcal{H}_c} \tilde{p}(z)\tilde{p}_0(z)}{\sum_{z \in \mathcal{H}_c} \tilde{p}_0(z)^2} - 1 \right). \quad (6.22)$$

Here $B(t)$ and $B_0(t)$ are the total probabilities of being in the subspace \mathcal{H}_c at time t in noisy and clean evolution, respectively, and \tilde{p} and \tilde{p}_0 are the corresponding bitstring probabilities normalized in \mathcal{H}_c . We numerically confirm that Eq. (6.22) is a good approximation in the strong Rydberg blockade regime, provided that noisy evolution results in negligible many-body overlap in the manifold outside \mathcal{H}_c . In addition, we take into account the spatial inversion symmetry of the Rydberg Hamiltonian. Specifically, our Hamiltonian commutes with the left-right inversion operator \hat{Q} , which swaps two atoms at sites i and $N - i + 1$ for every i , due to the global uniformity implicit in the one-dimensional Rydberg model. Under this symmetry, we find that $|z_e\rangle = \frac{|z\rangle + |\bar{z}\rangle}{\sqrt{2}}$ and $|z_o\rangle = \frac{|z\rangle - |\bar{z}\rangle}{\sqrt{2}}$ are the even- and odd-parity eigenstates of \hat{Q} with eigenvalues of 1 and -1 , respectively. Here, $|z\rangle$ and $|\bar{z}\rangle$ are the z -basis bitstring and its mirrored version, i.e., $|z\rangle = |z_1 z_2 \dots z_N\rangle$ and $|\bar{z}\rangle = |z_N z_{N-1} \dots z_1\rangle$ where z_i is either 0 or 1 at site i . Since our initial state $|\psi_0\rangle = |0\rangle^{\otimes N}$ is the even-parity eigenstate of \hat{Q} , i.e., $\hat{Q}|\psi_0\rangle = +|\psi_0\rangle$, the resulting many-body states after a quench always reside within the even-parity sector, reducing the effective Hilbert space dimension approximately by a factor of two. Furthermore, since projective measurement yields either $|z\rangle$ or $|\bar{z}\rangle$ with probabilities 1/2 due to the parity eigenbasis $\{|z_e\rangle, |z_o\rangle\}$, their contributions to the fidelity estimation formula need to be adjusted; in consideration of this inversion symmetry, we rewrite both $\tilde{p}(z)$ and $\tilde{p}_0(z)$ using the parity eigenbasis and assume that both $|z\rangle$ and $|\bar{z}\rangle$ measurement outcomes originate from the same $|z_e\rangle$ in the even-parity sector. This effectively corresponds to checking the bitstring correlation in the even-parity sector of the parity eigenbasis. Having incorporated these modifications, we numerically confirm that the adjusted fidelity estimator shows good agreement with the true fidelity defined from the *full* Hilbert space.

6.2.6 Errors in fidelity benchmarking

Fidelity estimation is not without error; besides the systematic effect of the delay time previously discussed, there are also statistical limitations. These originate from two sources: Sampling (i.e. we need to take enough samples for our benchmarking to be statistically significant) and Typicality (i.e. we need the Hilbert space to be large enough that some of our approximations in deriving F_c are always true). We note that we repeat these studies in perhaps a more intuitive way in Section 8.6.3, where we also show more experimentally relevant values for these error sources, but we present information on them here as well for completeness.

Sampling

We quantify the typical statistical error from approximating our fidelity estimator via Eq. 6.9 in two steps. First, we use our *ab initio* error model to simulate the quantum evolution of the Rydberg Hamiltonian for system sizes from $N = 10$ to 22, from which we can calculate the exact value of F_c compared to error-free numerics. We then sample a finite number of M samples from the probability distributions produced from the error model simulation, apply Eq. 6.9, and plot the standard deviation of F_c as a function of M (Fig. 6.6a). We see a characteristic scaling of $\sigma(F_c) = A/\sqrt{M}$, where A is the *sample complexity*, expected to scale exponentially with N , and σ denotes the standard deviation. We perform a similar process directly on our experimental data by repeatedly subsampling the experimentally measured bitstrings in order to estimate the scaling of the standard deviation. We plot both error model and experimental results in Fig. 6.6b as a function of N , where the time is set as the N -dependent entanglement saturation time. By fitting the experimental (error model) results, we find $A \approx a^N$, with $a = 1.037(2)$ ($a = 1.039(2)$).

Typicality

Benchmarking results for RUCs are averaged over an ensemble of RUCs realized via sampling of two-qubit unitaries \hat{U} . For F_{XEB} , a single choice of an RUC out of such an ensemble already yields accurate results in the sense that, in the limit of large circuit depth, the error from using only a single RUC scales as $1/\sqrt{D}$ with D being the Hilbert space dimension of a global system [162].

We now investigate if a similar typicality scaling holds for benchmarking of Hamiltonian dynamics. To this end, we consider an ensemble of Hamiltonians generated from a mixed-field Ising model [322], given as

$$\hat{H}_{\text{QIMF}}/\hbar = 2\pi \left[\sum_{i=1}^N 0.22\hat{S}_i^x + 0.25\hat{S}_i^y + \sum_{i=1}^{N-1} \hat{S}_i^x \hat{S}_{i+1}^x \right], \quad (6.23)$$

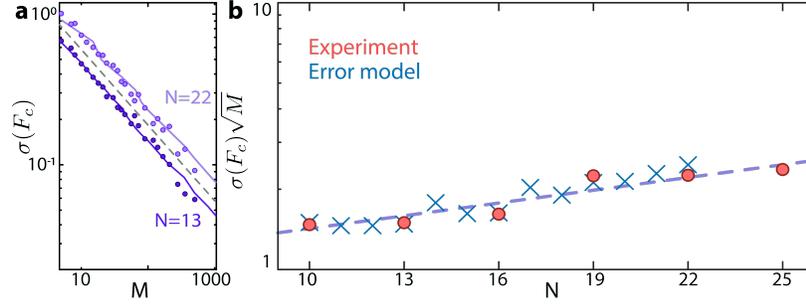


Figure 6.6: **Finite sampling analysis for F_c .** **a**, Statistical fluctuations of the fidelity estimator, F_c , at $N=13$ (dark purple) and $N=22$ (light purple), computed both using our *ab initio* error model (solid lines) and experiment (markers) evaluated with a finite number of M bitstring samples. Data are consistent with a $1/\sqrt{M}$ scaling, shown here as a guide to the eye (grey dashed line) **b**, Sample complexity of the fidelity estimator, evaluated at the N -dependent entanglement saturation time for the error model (blue crosses), and for the experimental data in Fig. 6.3d (red circles). A fit to the experimental data (dashed line) with functional form $\sigma(F_c)\sqrt{M} = a^N$ yields an estimate of $a = 1.037(2)$ (a similar fit to the error model yields an estimate of $a = 1.039(2)$).

by adding a series of site-dependent transverse fields:

$$\hat{H}_{\text{QIMF}}(\mathbf{J})/\hbar = \hat{H}_{\text{QIMF}}/\hbar + 2\pi \sum_{i=1}^N J_i \hat{S}_i^z, \quad (6.24)$$

where $\hat{H}_{\text{QIMF}}(\mathbf{J})$ is a Hamiltonian in the ensemble parameterized by site-dependent transverse fields, $\mathbf{J} = (J_1, \dots, J_i, \dots, J_N)$. We sample the J_i uniformly from $[-0.5, 0.5]$ with a restriction, $\sum_{i=1}^N J_i = 0$, imposed so all target states thermalize to effective infinite effective temperature (see Appendix C). With the $\hat{H}_{\text{QIMF}}(\mathbf{J})$, we generate an ensemble of parameterized target states, $|\psi(t, \mathbf{J})\rangle = e^{-i\hat{H}_{\text{QIMF}}(\mathbf{J})t/\hbar}|0\rangle^{\otimes N}$. We now consider a single error occurring at a time t_0 , induced by a single-site rotation around the x -axis, on top of otherwise error-free evolution. We denote the erroneous states as $|\psi(t, \mathbf{J})\rangle_e$. The rotation angle is fixed across \mathbf{J} -choices with an amplitude that yields a many-body overlap $F(\mathbf{J}) = |\langle \psi(t, \mathbf{J}) | \psi(t, \mathbf{J}) \rangle_e|^2 \approx 0.5$.

We now estimate this many-body overlap using the bitstring probabilities for these two states, yielding a time- and \mathbf{J} -dependent estimator $F_c(t, \mathbf{J})$. To quantify relative fluctuations, we introduce the relative difference $d(\tau, \mathbf{J}) = (F_c(t_0 + \tau, \mathbf{J}) - F(\mathbf{J}))/F(\mathbf{J})$, where $\tau = t - t_0$ is the evolution time after the error. We show the ensemble average of $d(\tau, \mathbf{J})$ as a function of τ (Fig. 6.7a) for two different choices of t_0 : $t_0 = 0$ and $t_0 = 5$. We observe that the ensemble average settles quickly to zero within a 10^{-2} level; however, we find fluctuations around this mean value arising from different choices of \mathbf{J} . To quantify these fluctuations,

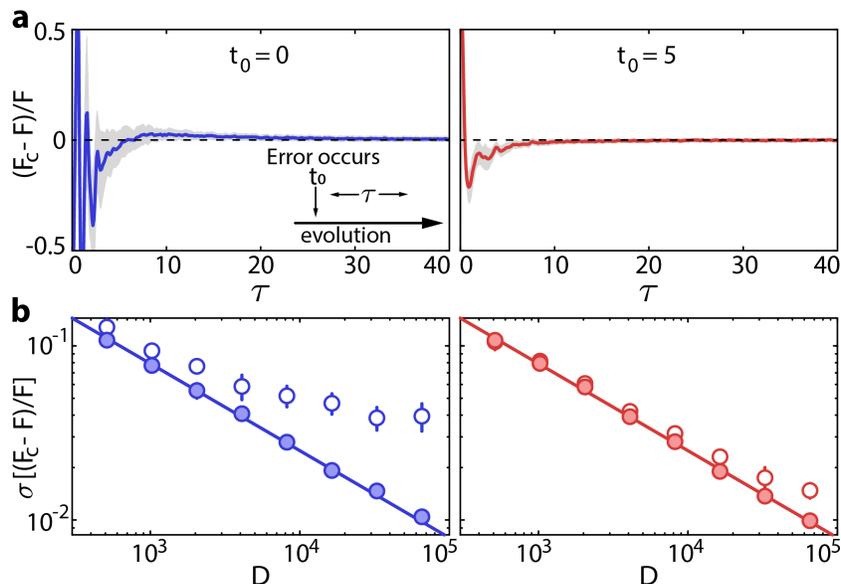


Figure 6.7: **Typicality analysis of F_c** : **a**, Normalized fidelity estimation errors as a function of interrogation time τ after a local, instantaneous error occurs at time t_0 (see the inset of **a**). The error strength is chosen such that the many-body overlap is reduced to ~ 0.5 . To investigate typical statistical errors in fidelity estimation, we simulate an ensemble of chaotic Hamiltonians at a fixed system size of $N = 16$ for two error occurrence times of $t_0 = 0$ (left) and $t_0 = 5$ (right). The mean and standard deviation of the normalized fidelity estimation errors are plotted as the solid line and grey band, respectively. **b**, Standard deviation of normalized fidelity estimation errors, characterized at two different post-error interrogation times ($\tau = 8$; open markers, $\tau = 30$; closed markers) as a function of the total Hilbert space dimension D for the two error occurrence times of $t_0 = 0$ (left) and $t_0 = 5$ (right). At the late interrogation time of $\tau = 30$, we find that the typical error from the distinct choice of a Hamiltonian follows a scaling of $\sigma[(F_c - F)/F] \sim 1/\sqrt{D}$ (solid lines), as also found for deep random unitary circuits. In **b**, error bars denote the standard deviation of $\sigma[(F_c - F)/F]$ in a $[\tau - 1, \tau + 1]$ window.

we evaluate the standard deviation of $d(\tau, \mathbf{J})$ over \mathbf{J} and find that it decreases with system size (Fig. 6.7b). At late times τ , we find a $1/\sqrt{D}$ scaling similar to the case of RUCs. These results suggest that, at sufficiently late time, our method becomes increasingly more precise in the limit of large system size for a single typical Hamiltonian evolution.

6.2.7 Benchmarking with evolution constraints

The fidelity estimator considered in the last section, F_c , accurately tracks evolution fidelity when that evolution is tuned such that the initial state thermalizes at infinite effective temperature (see Appendix C). However, problems may arise when this infinite effective temperature condition is not met, or if symmetries of the Hamiltonian restrict the system

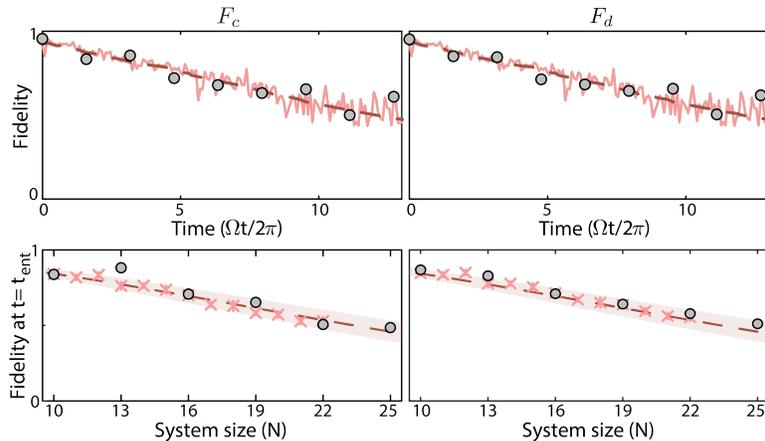


Figure 6.8: **Comparison to a new fidelity estimator.** We compare benchmarking with the fidelity estimator F_c , as in Fig. 6.3 (left), against a new fidelity estimator, F_d [30], which normalizes all probabilities in Eq. 6.8 by their time-average (right). Colors and marker types are equivalent to those in Fig. 6.3 in all subplots.

dynamics, as the projected ensemble may not completely form a 2-design. In recent work [30], a new estimator, F_d , was introduced by rescaling the bitstring probabilities in Eq. 6.8 by their time-average; remarkably this allows for benchmarking systems evolving at finite effective temperature as well as systems of itinerant particles such as Bose- and Fermi-Hubbard models. In Fig. 6.8 we compare this new estimator, F_d [30], and our original estimator F_c for benchmarking our Rydberg quantum simulator. We observe no sizeable difference between the two estimators, which we interpret as due to intentionally performing all evolution close to infinite effective temperature in this work. We revisit this improved estimator in detail in Chapter 8.

6.3 APPLICATIONS

Our study of random state ensemble led us to successfully benchmark the many-body fidelity for performing time-independent Hamiltonian quench dynamics. While already an important milestone that will be (and already is) useful for evaluating and improving quantum devices, there are several other more direct and practical application of our techniques. In this chapter we delve into such applications in more detail, highlighting three in particular; a fourth application, namely learning microscopic noise models affecting our system will be discussed in further detail in the following Chapter.

First, we perform *Hamiltonian parameter estimation*, or Hamiltonian learning, via many-body fidelity measurements. Specifically, we perform a many-body quench with the ex-

periment, and then in the corresponding simulation we vary the Hamiltonian parameters. We find that a clear maxima emerges which agrees with pre-calibrated values, and which gives access to space- and time-dependent parameters, and for instance may be useful for uncovering unexpected Hamiltonian terms.

Second, we extend the reach of our fidelity estimation from simple quench dynamics to a broader class of *target state benchmarking*. Specifically, we numerically show that the fidelity for many types of complex quantum processes can be benchmarked by appending a short chaotic quench after the dynamics, and measuring the result fidelity. We numerically find this technique works for adiabatic sweeps as well as for specific interesting entangled states. This task is spiritually quite related to Hamiltonian parameter estimation, as both are essentially many-body fidelity estimation with different choices of target or initial states.

Finally, we demonstrate that many-body fidelity estimation allows for evaluating the relative performance of analog and digital quantum devices, even when their dynamics are quite different. In this case we will use a relatively simple metric, which is the effective gate fidelity per unit entanglement time, but we note that in Chapter 9 we greatly extend this formalism to capture the actual mixed state entanglement content of experiments.

6.3.1 Hamiltonian parameter estimation

First, we demonstrate coherent Hamiltonian parameter estimation. To do so we measure F_c while varying Hamiltonian parameters in simulation; when numerical parameters do not match with those in experiment, F_c will decay quickly. To capture this effect in a single quantity we plot the normalized, time-integrated F_c (Fig. 6.9a). For each Hamiltonian parameter, a sharp maximum emerges, showing good agreement with precalibrated values (dashed lines and shaded areas). Parameter estimation also works when applied to learn local, site-dependent terms of a disordered Hamiltonian (Fig. 6.9b), notably without any local control during readout.

The Hamiltonian parameter estimation scheme presented in Fig. 6.9 works by comparing measurements from experiment against an ensemble of target states from numerical simulations, each evolved with a different set of Hamiltonian parameters. For example, we can define a family of target states, which are parameterized by the Rabi frequency, Ω , as $|\psi(t, \Omega)\rangle = e^{-it\hat{H}(\Omega)/\hbar}|0\rangle^{\otimes N}$. When the value of Ω does not match the Rabi frequency used in the experiment, the target state $|\psi(t, \Omega)\rangle$ will have smaller overlap with the experimental state, and the fidelity estimator $F_c(t, \Omega) \approx \langle \psi(t, \Omega) | \hat{\rho}(t) | \psi(t, \Omega) \rangle$ will decay more quickly. To identify the best Rabi frequency yielding the highest fidelity, we define the

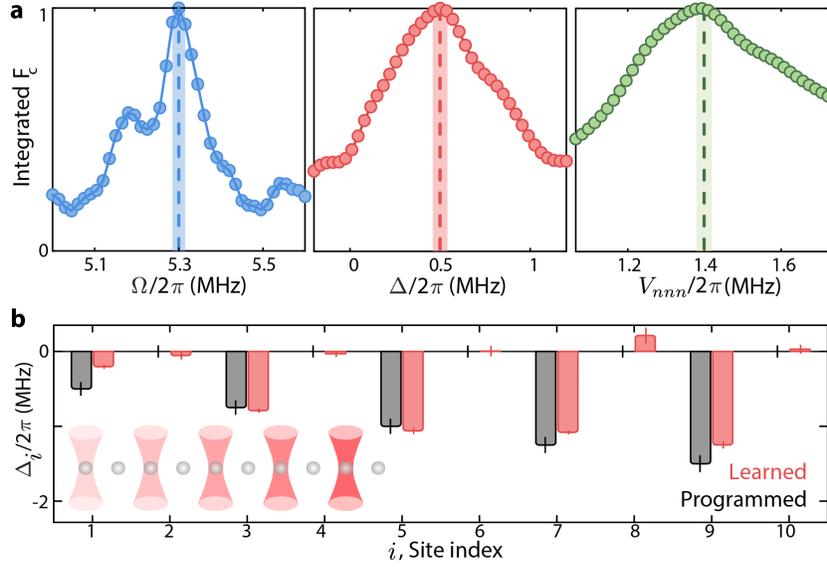


Figure 6.9: **Hamiltonian parameter estimation.** **a**, Normalized, time-integrated F_c as a function of the global Rabi frequency, detuning, and the next-nearest-neighbor interaction strength in the Rydberg model; this normalized F_c is maximized only when the correct parameters are used. Vertical dashed lines and shaded areas denote independently calibrated values and their uncertainties. **b**, Programmed (grey bars) and learned (red bars) local Hamiltonian parameters for an arbitrary, site-dependent detuning field imposed with an intensity-dependent lightshift from locally addressable optical tweezers (inset, red funnels).

time-integrated estimator defined as $\bar{F}_c(\Omega) = \int_0^T dt F_c(t, \Omega)$, where T is the maximum interrogation time, and plot $\bar{F}_c(\Omega)$ as a function of Ω .

The Hamiltonian parameter estimation scheme presented in Fig. 6.9 allows for estimating both global and local Hamiltonian parameters; here we show how its sensitivity compares to, for instance, a metric based on the residual-sum-of-squares between the local magnetization from experiment (or any noisy reference) and from error free numerics,

$$RSS(\Omega, t) = \sum_{i=1}^N (\langle \hat{S}_z^i(\Omega) \rangle - \langle \hat{S}_{z,\text{exp}}^i \rangle)^2, \quad (6.25)$$

where \hat{S}_z^i the local magnetization along the z axis for the atom at site i from the numerical simulation, and $\hat{S}_{z,\text{exp}}^i$ is the same magnetization from a fixed reference, for instance the experimental data. Here, Ω can be regarded as *any* parameter that determines the Hamiltonian of interest (e.g. Rabi frequency). We hypothesize that while the RSS will be peaked at the correct set of Hamiltonian parameters, at late times it will be less sensitive than our proposed protocol based on F_c because while the local magnetization will thermalize to a fixed value, the global state, which F_c is sensitive to, will still be continuously evolving.

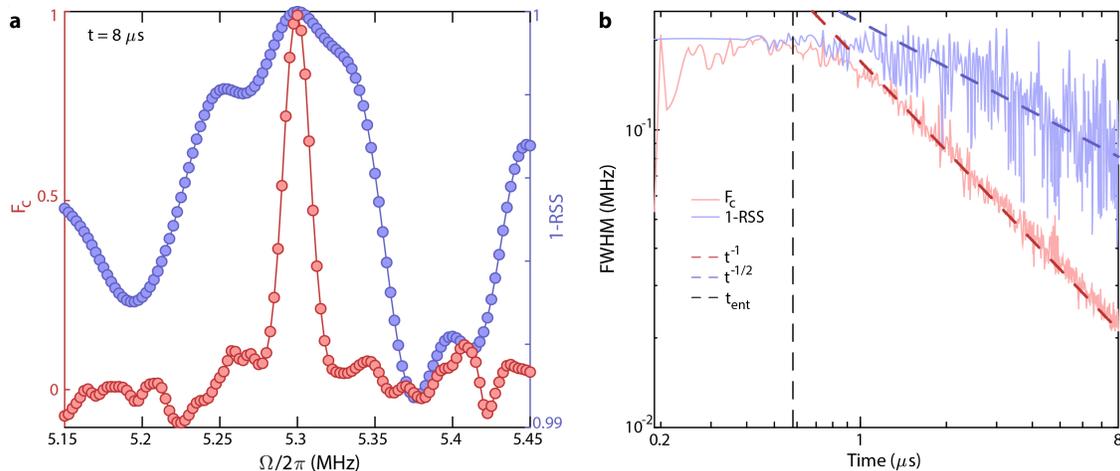


Figure 6.10: Sensitivity of Hamiltonian parameter estimation. **a**, We perform an error-free numerical simulation of the Rydberg Hamiltonian with $N = 12$ and fixed Rabi frequency, $\Omega = 5.3$ MHz, which is then compared against another error-free simulation with varying Ω . To explore the efficacy of our parameter estimation scheme, we compare the two simulations in two ways: via F_c (red, as in Fig. 5a of the main text), and via $1 - RSS$, where RSS is the residual-sum-of-squares of the local magnetizations (blue, see text); both estimators are taken at a single time point of $t = 8 \mu\text{s}$. We note that F_c is both more sharply peaked at the correct Ω , and that the range of values which F_c takes is larger than for $1 - RSS$, which has a dynamic range of only 0.01. **b**, The FWHM of the central peak in F_c in **a** stays relatively flat up until t_{ent} , the entanglement time of the system (see Fig. 4d of the main text for example), before decreasing as $1/t$ at late times. The FWHM of the RSS estimator follow similar behavior, but follows a $1/\sqrt{t}$ trend at late times.

We test this hypothesis via error-free numerical simulation: first we simulate the Rydberg Hamiltonian with the Rabi frequency fixed at $\Omega/2\pi = 5.3$ MHz – this set of numerics serves as the reference. Next, we repeat the simulation several times, each time with a different Rabi frequency, varying from $\Omega/2\pi = 5.15$ MHz to 5.45 MHz. We then compare these simulations on the level of F_c and RSS to identify how well both estimation methods can identify 5.3 MHz as the correct Rabi frequency in the original simulation.

Such results are plotted in Fig. 6.10a at a late fixed time of $t = 8 \mu\text{s}$ ($\Omega t/2\pi \approx 42$). We see that indeed while both F_c and $1 - RSS$ are peaked at $\Omega/2\pi = 5.3$ MHz, the F_c peak is both more sharp, and has a larger dynamic range. Further, we plot the central peak width as a function of time in Fig. 6.10b, where we see after the entangling time, t_{ent} (given as the point at which the half-chain entanglement entropy saturates), the FWHM of the F_c central peak decreases as $\approx 1/t$, while the FWHM of the $1 - RSS$ central peak decreases as $\approx 1/\sqrt{t}$, indicating that the protocol based on F_c will become comparatively more sensitive as the quench time increases.

6.3.2 Target state benchmarking

Next, F_c can be used to benchmark the fidelity of preparing various quantum states of interest by preparing a target state and then quenching the Hamiltonian to evolve the prepared state at infinite effective temperature (Fig. 6.11a). In this modality, we assume the target state is prepared with some non-unity fidelity due to experimental imperfections, after which we apply an infinite effective temperature quench Hamiltonian, and observe the resulting dynamics.

As a numerical proof-of-principle, we show results for such target state benchmarking to prepare a ground state near the Ising quantum phase transition in the Rydberg model [323] (Fig. 6.11a,b), where the noisy state is an equal probability mixture of the ground and first excited states. After a short disordered quench, the estimator F_c reveals the fidelity of the prepared state, offering a novel way to perform *in situ* optimization of many-body state preparation.

In more detail, in Fig. 6.11a, the ideal state is the ground state at $\Delta/\Omega = 3$, $V_{\text{nnn}}/\Omega = 0.26$, close to the phase transition between the disordered and \mathbb{Z}_2 -ordered states of the Rydberg Hamiltonian [323]. The imperfect state is taken to be an incoherent mixture composed of 50% each of the ground and first excited states. This state is then quenched with a Hamiltonian with parameters $\Omega/2\pi = 5.3$ MHz, $\Delta/2\pi = 2.8$ MHz, $C_6/2\pi = 254$ GHz μm^6 , $a = 3.75$ μm , with $2\pi \times \pm 1$ MHz random on-site disorder drawn from a uniform distribution.

In Fig. 6.11 we provide further examples of the technique by numerically demonstrating fidelity estimation of various target states such as a cluster state, a Haar-random state of a 2D Rydberg quantum simulator, and a symmetry-protected topological (SPT) ground state.

Specifically, in Fig. 6.11c, we estimate the state preparation fidelity of a one-dimensional cluster state defined as

$$|\psi\rangle_{\text{cluster}} = \prod_{i=1}^{N-1} (\text{CZ})_{i,i+1} |+\rangle^{\otimes N} \quad (6.26)$$

where $(\text{CZ})_{i,i+1}$ is the two-qubit, controlled- Z gate acting on two adjacent qubits at site i and $i + 1$, and $|+\rangle$ is the equal superposition of the $|0\rangle$ and $|1\rangle$ states. The imperfect quantum state is prepared by applying a global phase rotation to the ideal state such that the state overlap becomes $F = 0.5$. We then employ an infinite effective temperature quench Hamiltonian given as $\hat{H}/\hbar = h_x \sum_i (\hat{S}_i^x - 1.79\hat{S}_i^y + 4.64\hat{S}_i^x \hat{S}_{i+1}^x)$ to learn the state overlap via our F_c formula.

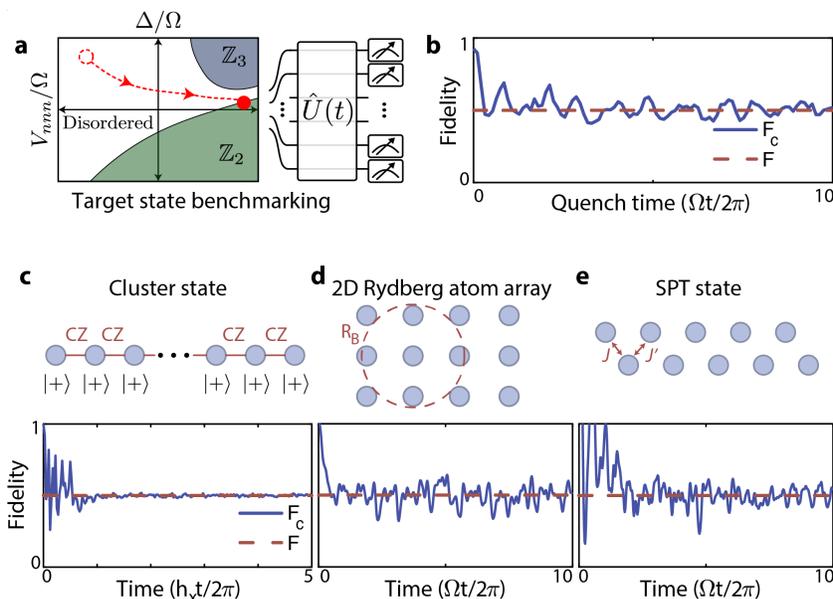


Figure 6.11: **Applications to target state benchmarking.** **a,b**, Our F_c benchmark (solid line) can estimate the fidelity F (dashed line) of producing a specified target state by evolving at infinite effective temperature after preparation, here numerically demonstrated for a ground state of system size $N = 15$ near the \mathbb{Z}_2 Ising quantum phase transition in the one-dimensional Rydberg ground state phase diagram [323], with a noisy state consisting of an equal probability mixture of the ground and first excited states. **c**, Benchmarking of a one-dimensional cluster state, **d**, a pure Haar-random state benchmarked in a two-dimensional square Rydberg atom array, and **e**, a symmetry-protected topological (SPT) ground state prepared in a Rydberg ladder array realizing the Su-Schrieffer-Heeger topological model [44]. In **c**, CZ denotes a controlled- Z gate and $|+\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$. In **d**, R_B denotes the Rydberg blockade radius within which more than a single excitation is not allowed [34, 36, 41]. In **e**, J and J' are the alternating coupling strengths of a two-leg ladder array, respectively. In all cases, $N = 16$ qubits are used, and imperfect quantum states are prepared via phase rotations such that the many-body fidelity overlap becomes 0.5 (red dashed line). Additionally, chaotic evolution is performed such that the initial state is at infinite effective temperature to apply to our F_c formalism (blue solid lines).

In Fig. 6.11d, we estimate the state preparation fidelity of a pure Haar-random state generated from Rydberg atoms in a 4×4 two-dimensional square array. The imperfect state is prepared by applying a local phase rotation to a central qubit, yielding $F = 0.5$. For subsequent quench dynamics, identical Hamiltonian parameters are used as in the 1D Rydberg benchmarking case (Fig. 6.3).

Lastly, in Fig. 6.11e, we estimate the state preparation fidelity of a symmetry-protected topological (SPT) ground state prepared in a Rydberg ladder array realizing the Su-Schrieffer-Heeger topological model, following the approach of Ref. [44]. The imperfect state is prepared with a local phase error yielding a state overlap of $F = 0.5$. It is subsequently

benchmarked via infinite effective temperature evolution with the quench Hamiltonian chosen to be the combination of the identical interaction Hamiltonian, a random on-site disorder of strength 1 MHz, and a detuned global drive with a Rabi frequency of 2 MHz and a detuning of 0.5 MHz.

6.3.3 Cross-platform comparisons

Finally, to compare analog and digital quantum evolution, we evaluate the fidelity achieved at t_{ent} for both analog quantum simulators and digital quantum computers (for which t_{ent} is defined in terms of gate depth). We find our system has an equivalent effective, SPAM-corrected, two-qubit cycle fidelity of 0.987(2) for the gate-set used in Ref. [162], and 0.9965(5) for a gate-set based on two-qubit SU(4) gates [277] (Fig. 6.12), and below provide details behind this calculation.

We note that while this estimation method gives a useful general methodology for comparing quantum systems of various types, in Chapter 9 we will discuss and implement an even more general formalism based on direct estimation of mixed state entanglement (rather than simply an entanglement-fidelity product, as we have implicitly utilize here).

Calculating the entanglement time for RUCs

For the case of finding the entanglement saturation depth, d_{ent} , for the case of random unitary circuits (RUCs) considered in Fig. 6.12, we apply essentially the same procedure. We study two different digital circuit implementations. In the first, with a gate-set based on Ref. [162], the RUC circuit is composed of alternating one- and two-qubit gates; the one-qubit gates are randomly chosen $\pi/2$ rotations along the \hat{x} , \hat{y} and $\hat{x} + \hat{y}$ directions, while the two-qubit gates are ‘fSym’ [162]. In the second, the RUC is composed entirely of two-qubit SU(4) gates (without global swap operations [277]). For the first implementation, we set $C = 1$, while in the second we set $C = 1.7$, to better guarantee the chosen depth is in the saturated entanglement regime. Open boundary conditions are used in accordance with the experimental Rydberg system, and thus there are two possible gate topologies (i.e. in the first depth applying the ‘fSym’ gate to qubits 1-2, 3-4, etc. or 2-3, 4-5, etc.) – we explicitly average over an equal number of randomized realizations of each topology when calculating the entanglement entropy growth. As in the Rydberg case, we fit the entanglement saturation depth as a linear function of the number of qubits in the RUC, yielding $d_{\text{ent}} = \beta_0 + \beta_1 N_{\text{RUC}}$. For the gate-set based on Ref. [162] we find $\beta_0 = -0.395(17)$, and

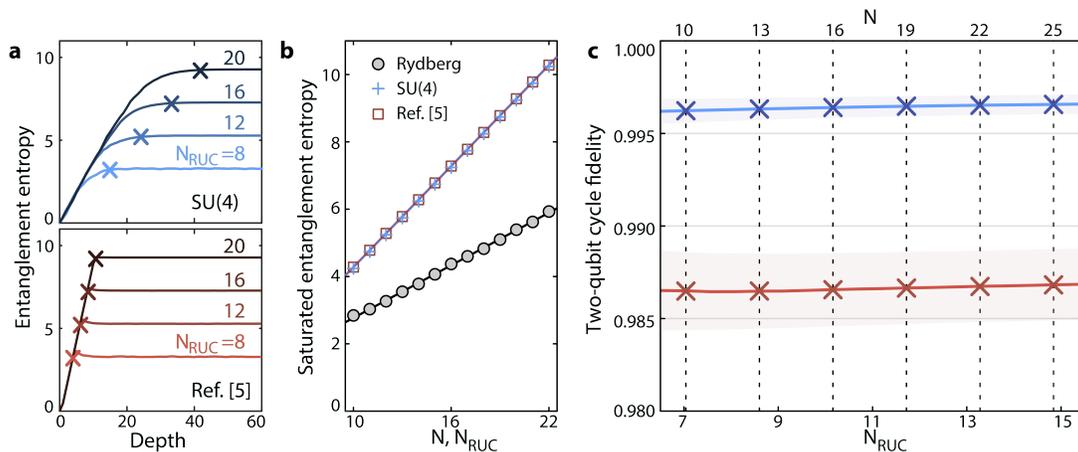


Figure 6.12: **Comparison to digital quantum devices executing random circuits.** **a**, Numerical simulations of a one-dimensional digital quantum device implementing a random unitary circuit (RUC). Two different digital gate implementations are tested: a configuration based on the gate-set used in Ref. [162] (bottom), and an configuration where each cycle is composed of parallel two-qubit SU(4) gates (top) [277]. Cross markers indicate when the half-chain entanglement entropy saturates. **b**, Due to the Rydberg blockade mechanism, as well as symmetries of the Rydberg Hamiltonian (Supplementary Information), an equal number of atoms in the Rydberg simulator, N , and qubits in the RUC, N_{RUC} , will not saturate to the same half-chain entanglement entropy. However, we can still find an equivalence by plotting the saturated entanglement entropy for the RUC (blue crosses for the SU(4) gate-set, open red squares for gate-set from Ref. [162]) and for the Rydberg simulator (grey markers) as a function of their respective system sizes. We fit the results for the Rydberg simulator (black line), and plot the analytic prediction for the RUC [324] (purple line), from which we can write an equivalent N_{RUC} as a function of N , in the sense of maximum achievable entanglement entropy (Methods). **c**, For a given N (and equivalent N_{RUC}), we plot the SPAM-corrected, two-qubit cycle fidelity for an equivalently-sized RUC to match the evolution fidelity of our Rydberg simulator at the time/depth when entanglement saturates. Red lines, markers and crosses are for the gate-set of Ref. [162], while blue are for the SU(4) gate-set. Shaded regions come from the error on fitting the various N -dependent parameters which enter this calculation (Methods).

$\beta_1 = 0.557(1)$, and for the gate-set based on SU(4) gates [277] we find $\beta_0 = -3.18(77)$, and $\beta_1 = 2.261(51)$.

Estimating fidelity decay

In Fig. 6.3, we see the decay profile of the model fidelity, F_{model} , for our Rydberg simulator is approximately exponential. We confirm this result roughly in our error model (though shall return to this point in Chapter 8.4). We perform quench evolution in the error model

for various system sizes (not shown) and fit the fidelity decay profile as,

$$F(N, t) \propto \exp(-\gamma(N)t), \quad (6.27)$$

where $\gamma(N)$ is the fidelity decay rate. We find that for the system size range considered here, $\gamma(N)$ scales approximately linearly with N , from which we fit $\gamma(N) = \gamma_0 + \gamma_1 N$; for our particular Hamiltonian parameters and noise sources, we find $\gamma_0 = 0.12(4)$, and $\gamma_1 = 0.017(3)$, both in MHz.

For the RUC case, fidelity decay for a given system size, N_{RUC} , and depth, d , is modeled as a simple product over constituent two-qubit cycle fidelity, F_{cycle} , yielding

$$F_{\text{RUC}}(N, d) = F_{\text{cycle}}^{(N_{\text{RUC}}-1)d/2}, \quad (6.28)$$

where the exponent of the right-hand side reflects the fact that we apply, on average, $(N_{\text{RUC}} - 1)/2$ two-qubit gates in parallel per depth.

Note that in Fig. 6.3, we used this fitted $\gamma(N)$ explicitly to predict the fidelity scaling of our Rydberg simulator at the N -dependent entanglement saturation time, t_{ent} , as a function of system size. Concretely, we plotted (red dashed line):

$$F_{\text{model}}(N, t_{\text{ent}}(N)) = F_0^N \exp(-\gamma(N)t_{\text{ent}}(N)), \quad (6.29)$$

where $F_0=0.997(1)$ is the single-atom preparation fidelity determined experimentally. The shaded red region in Fig. 6.3 depicts the error from fit uncertainty of $\gamma(N)$.

Comparing digital and analog devices

We wish to directly compare the evolution fidelity of our analog Rydberg quantum simulator against that of a digital device implementing an RUC with equivalent entanglement entropy at the entanglement saturation time. However, due to the Rydberg blockade mechanism, as well as symmetries of our Hamiltonian (Supplementary Information), an equal number of atoms in the Rydberg simulator, N , and qubits in the RUC, N_{RUC} , will not saturate to the same half-chain entanglement entropy.

To overcome this, in Fig. 6.12b we plot the entanglement entropy, S (S_{RUC}), achieved at t_{ent} (d_{ent}) for the Rydberg simulator (RUC) as a function of N (N_{RUC}). For the Rydberg simulator, we fit $S(N) = \sigma_0 + \sigma_1 N$ with $\sigma_0 = 0.16(4)$, $\sigma_1 = 0.26(3)$. For the RUC, we use the prediction of $S_{\text{RUC}}(N_{\text{RUC}}) = \eta_0 + \eta_1 N_{\text{RUC}}$ with $\eta_0 = -\log_2(e)/2 \approx -0.72$ and $\eta_1 = 1/2$ being exact values with no error bars, as defined in Ref. [324] (where e is Euler's number, and where we have used the \log_2 entanglement entropy convention). To find the

equivalent N_{RUC} for a given N , we then simply equate $S_{\text{RUC}}(N_{\text{RUC}}) = S(N)$, yielding $N_{\text{RUC}} = (\sigma_1 N + (\sigma_0 - \eta_0)) / \eta_1 = 0.52N + 1.76$.

With this system size equivalence established we can directly compare the SPAM-corrected Rydberg and RUC systems at their respective entanglement time and depth, in order to find the equivalent RUC two-qubit cycle fidelity which would match the Rydberg quantum simulator's evolution fidelity. By evolution fidelity, we refer to the fidelity at the entanglement time, up to preparation errors, which based on our validated error model is approximately given by $\exp(-\gamma(N)t_{\text{ent}}(N))$ from Eq. 6.29. We equate $F_{\text{RUC}}(N_{\text{RUC}}, d_{\text{ent}}) = \exp(-\gamma(N)t_{\text{ent}}(N))$, and then solve for F_{cycle} . As shown in Fig. 6.12c, for the gate-set used in Ref. [162], we find $F_{\text{cycle}} = 0.987(2)$, while for the SU(4) circuit we find $F_{\text{cycle}} = 0.9965(5)$, nearly independent of system size. Error bars originate from the uncertainty on the parameters of γ , t_{ent} , d_{ent} , F_0 , and S .

Hilbert-space ergodicity

The developing study of *quantum thermalization* [305, 307, 308, 325–329] has sought to create a consistent formulation for both statistical and quantum mechanics. In this framework, entanglement entropy between subsystems of a quantum state plays the role of thermal entropy in statistical mechanics. Considering a closed, bipartite quantum system composed of interacting subsystems A and B under a time-independent Hamiltonian (Fig. 7.1a), quantum thermalization suggests that at late times, local observables in A will behave as if they were drawn from a thermal state. Here B acts as a thermal bath, for which the system Hamiltonian and initial state define an effective temperature [329], as supported by numerous experiments [29, 306, 330–333]. See Appendix C for a more pedagogical introduction to the notion of effective temperature.

Quantum thermalization predicts expectation values of *local* observables, but newer theoretical perspectives have begun to shed light on the behavior of *global* quantum systems. In particular, it was proposed that global quantum states homogeneously explore their available Hilbert space [71, 334] (Fig. 7.1b). This phenomenon, dubbed *Hilbert-space ergodicity* (HSE), is analogous to ergodicity in classical mechanics, whereby over time a system visits its state space uniformly. The assumption of HSE entails consequences beyond standard predictions of quantum thermalization. For instance, consider a global projective observable which is not described by quantum thermalization, like the probability to measure a specific basis state (e.g. a ‘bitstring’ for a system of qubits). HSE predicts that *relative* probability fluctuations over time precisely follow an exponential (or Porter-Thomas [275, 289]) distribution, both for the well-known case of randomized circuits [275, 297], and even for time-independent Hamiltonian dynamics typically studied by quantum thermalization.

Here we observe both existing predictions of HSE and new generalizations we introduce,

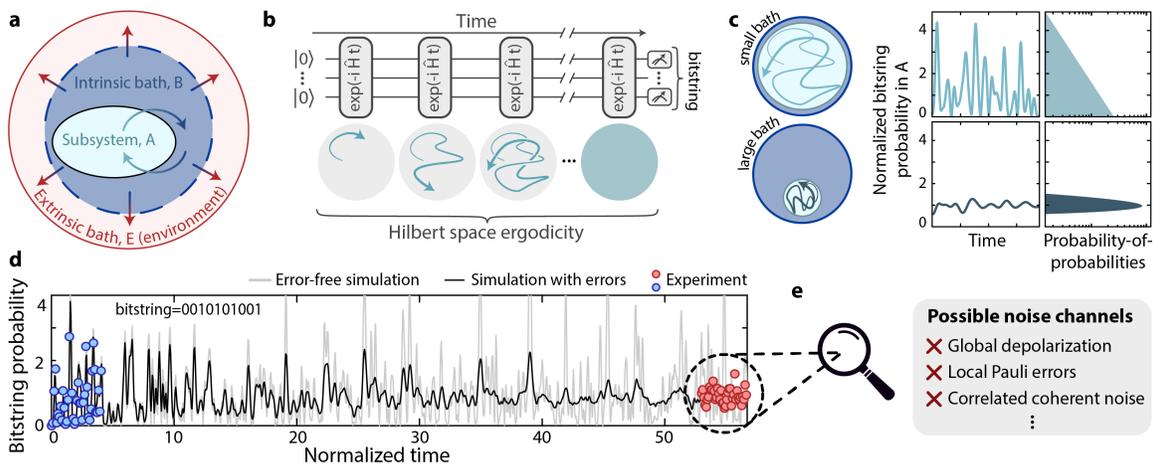


Figure 7.1: Quantum systems interacting with baths. **a.** A quantum system, A , interacts with an intrinsic bath, B , and with the external environment, E . **b.** We study temporal fluctuations of observables in A , and find they take on a universal form for closed, time-independent Hamiltonian systems exhibiting *Hilbert-space ergodicity* [30, 71], a recent theoretical notion that isolated quantum systems ergodically explore their available Hilbert space. Blue arrows are pure state trajectories. We verify these predictions on an experimental Rydberg quantum simulator, and numerically with various systems. **c.** For closed dynamics (i.e. with no coupling to E), we find an exact and universal dependence of the fluctuations of *normalized* bitstring probabilities in A on the intrinsic bath dimension: when the bath is small, relative fluctuations are large, and vice versa. This results in a universal probability-of-probabilities (PoP) distribution aggregated over measurement times. **d.** In the latter half of this work we study interactions with the external environment. As an example, we show bitstring probabilities following quench dynamics from experiment (markers), error-free simulation (grey line), and error-model simulation [31] (black line). For both the experiment and the error-model, probabilities become less fluctuating over time as errors accumulate, analogous to the closed system case where fluctuations decreased with increasing intrinsic bath dimension. **e.** Based on this heuristic connection, we introduce a framework to identify and discriminate between different types of system-environment coupling.

using an experimental Rydberg quantum simulator [31, 36, 41] and numerical models of spins, itinerant particles, and random circuits. Specifically, we observe the global Porter-Thomas distribution, and extend HSE to the case of a local system interacting with its environment to make concrete predictions beyond the scope of conventional quantum thermalization.

First, we consider the typical setting of quantum thermalization: when the environment, B , is a complementary subsystem of a closed system, which we call the intrinsic bath (Fig. 7.1a). For thermalization at infinite effective temperature, we find temporal fluctuations of projective observables follow a universal form: the Erlang distribution [291]. Solely parameterized by the dimension of the bath, the Erlang distribution undergoes a smooth

quantum-to-classical transition, progressing from the global Porter-Thomas distribution to a narrow Gaussian distribution as the intrinsic bath size increases (Fig. 7.1c). These results are exact and universal (i.e. independent of system details) for ergodic systems both globally and locally. This in contrast to studies of quantum thermalization which typically either only bounded [335, 336] or predicted the scaling of fluctuations with the total system dimension [337–339]; previous predictions have been made more exact for special cases like explicitly randomized dynamics [340] or the weakly non-integrable regime [341].

Then, we consider the case when the environment, E , is external, and the global composite system, AB , evolves under open system dynamics (Fig. 7.1a). Such a setting is omnipresent for quantum experiments [50]. The coupling between E and AB damps out fluctuations of observables like bitstring probabilities (Fig. 7.1d), but in a way which depends on the microscopic system-environment interaction. Based on our understanding of HSE principles, we devise a formalism to account for this behavior. We identify classes of noise models with distinct signatures on observable fluctuations in order to discriminate candidate models directly with many-body bitstring measurements (Fig. 7.1e). Our method can handle bespoke models with correlated noise and with heavy tails [342], both of which could be detrimental to quantum error correction [342–344].

Overall, our work studies the role of ergodicity in open and closed quantum dynamics, with fundamental and practical consequences for modern quantum experiments.

7.1 CLOSED SYSTEM DYNAMICS

First, we consider the case with no external environment (closed system dynamics), and examine the microscopic interactions of the subsystem, A , and its intrinsic bath, B , through the lens of HSE. The statement of HSE is that global quantum states uniformly cover their *available* Hilbert space, as constrained by conservation laws, symmetries, or a finite effective temperature. This is guaranteed when there are no high-order degeneracies in the energy spectrum, referred to as the no-resonance conditions (Section 7.4.2). When these conditions are fulfilled, at any particular (late) time we can consider the quantum state as essentially being sampled randomly from the ensemble of all possible states in the constrained Hilbert space. For instance, in the case of time-independent Hamiltonian evolution, this ensemble is the random phase ensemble [30, 71, 345], the set of quantum states with well-defined amplitudes in the energy eigenbasis (fixed by the initial state), but random phases.

This innate randomness then enables us to make quantitative predictions for the *relative*

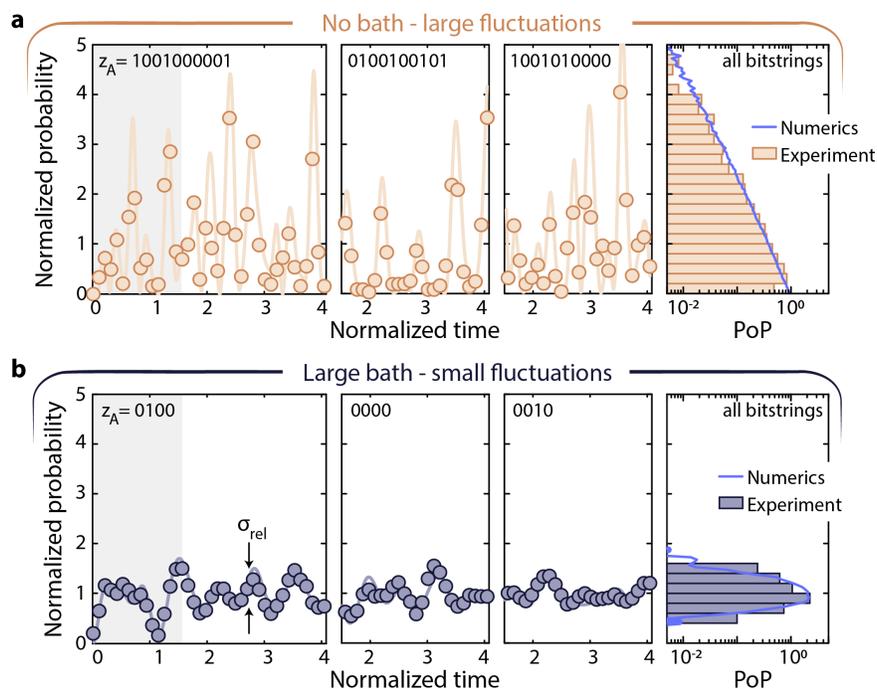


Figure 7.2: **Subsystem fluctuations are tied to the bath dimension.** **a.** When there is no intrinsic bath ($N_A = 10$ qubits, $N_B = 0$ qubits), normalized bitstring probabilities from both experiment and numerics are highly fluctuating past a short onset timescale (grey fill before this time, normalized time is in Rabi cycles), resulting in a wide, exponential PoP distribution (aggregated over all bitstrings and times). **b.** When the bath is large ($N_A = 4$ qubits, $N_B = 6$ qubits), fluctuations are substantially reduced (e.g. typified by the temporal variance, σ_{rel}^2), and the PoP distribution is narrow.

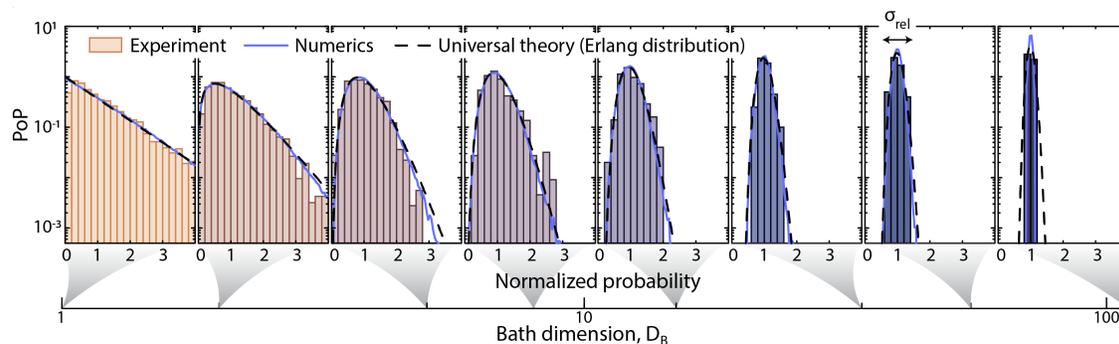


Figure 7.3: **Observing a smooth quantum-to-classical transition for closed system dynamics.** PoP distributions for many bath Hilbert-space dimensions, D_B , are in excellent agreement with our prediction that the PoP follows a universal form, the *Erlang distribution* (see text). The Erlang distribution starts as an exponential distribution for $D_B = 1$ before becoming progressively more narrow and Gaussian as D_B increases.

fluctuations of projective observables like the probability, $p(z_A, t)$, to measure a given bitstring in A , z_A (e.g. 1001), at a particular time, t . By relative, we mean probabilities are normalized by their long-time average, e.g. $\tilde{p}(z_A, t) \equiv p(z_A, t)/p_{\text{avg}}(z_A)$. Normalization by the mean eliminates the effects of symmetries and conservation laws which constrain HSE and make the $p_{\text{avg}}(z_A)$ non-equal [30, 71].

To study temporal fluctuations of \tilde{p} , we construct the probability-of-probabilities (PoP) distribution, a histogram, $P(\tilde{p})$, binning all \tilde{p} over time (Fig. 7.1c). The PoP encodes higher-order moments of \tilde{p} , e.g. the variance of the PoP is the variance of \tilde{p} over time. Remarkably, under HSE we can analytically derive the PoP (Section 7.4.3). In particular, for large global Hilbert-space dimension, $D \rightarrow \infty$, and for thermalization at infinite effective temperature, we predict the PoP to be

$$P_{\text{Erlang}}(\tilde{p}; D_B) = \frac{D_B^{D_B}}{(D_B - 1)!} \exp(-D_B \tilde{p}) \tilde{p}^{D_B - 1}, \quad (7.1)$$

which is the *Erlang distribution*, parameterized solely by the intrinsic bath Hilbert-space dimension, D_B .

This indicates that in the thermodynamic limit, temporal fluctuations of observables in subsystem A are entirely controlled by the dimension of the bath, and *not* by the dimension of the subsystem in which they are measured. For $D_B = 1$, the Erlang becomes the exponential (Porter-Thomas) distribution [275, 289], while for $D_B \rightarrow \infty$ it approaches [291] a narrow Gaussian distribution. We note previous theoretical studies [71] of HSE predicted the emergence of the Porter-Thomas distribution for global quantum systems, but not the more general Erlang distribution which is a novel theoretical discovery of our work. While this assumes an infinite effective temperature, we consider finite effective temperatures in Section 7.4.3.

We now test if experimental and numerical observations are consistent with our predictions. We implement a time-independent, global quench under the one-dimensional Ising-like Rydberg Hamiltonian [36, 41] with $N=10$ qubits. Parameters are chosen such that the initial all-zero state, $|\Psi_0\rangle = |0\rangle^{\otimes N}$, thermalizes at infinite effective temperature [308]. While our theoretical proofs are only rigorously true in the infinite time limit, for the experiment we focus on early time dynamics, defined below, such that the global fidelity remains high [31] (i.e. such that the influence of the external environment is small). Still, remarkably we find the key signatures of our theory are visible experimentally already at these early times.

Qubits are measured after evolving for a time t , producing a global bitstring $z \in \{0, 1\}^N$, and over many measurements we build up the global probability distribution $p(z, t)$. In post-

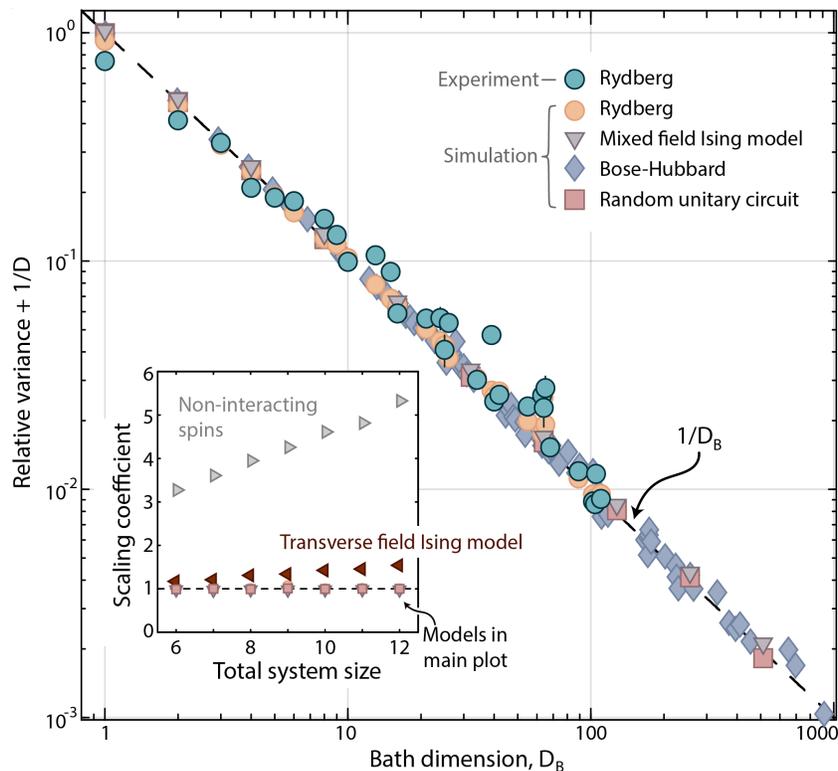


Figure 7.4: **Universality of the PoP variance.** The PoP variance (i.e. the variance of temporal fluctuations) scales inversely with the bath dimension for both the experiment and various numerical models exhibiting Hilbert-space ergodicity, in agreement with Eq. (7.2). Inset: The scaling coefficient $\langle D_B(\sigma_{\text{rel}}^2 + 1/D) \rangle$ ($\langle \cdot \rangle$ denotes an average over choices of D_B) is uniformly 1 for the models studied in the main plot, but deviates as a function of total system size for models for which Hilbert-space ergodicity does not hold. For the Bose-Hubbard system in the main plot, and all models in the inset, the bath dimension is rescaled to account for conservation laws and finite temperature, see Fig. 7.8.

processing, we bipartition into arbitrary subsystems A and B (of sizes N_A and N_B , and Hilbert-space dimensions D_A and D_B respectively), and study the corresponding marginal, normalized probability distributions $\tilde{p}(z_A, t)$. Normalization factors, $p_{\text{avg}}(z_A)$, are found numerically. We note that D_A and D_B are non-trivially modified due to the Rydberg blockade regime [36] in which we operate the experiment (Section 7.3).

First, we consider global probabilities (i.e. $N_A=N=10$ and thus $D_B=1$). We observe probabilities are highly fluctuating, and upon aggregating all probabilities from 1.5 to 4 cycles of evolution, we find the PoP is a nearly exponential distribution [30, 289], $P(\tilde{p}) = \exp(-\tilde{p})$, a concrete experimental observation consistent with HSE [71] (Fig. 7.2a). Then, for local systems (e.g. for $N_A = 4$ and $N_B = 6$), we find probability fluctuations are greatly diminished, and the PoP is accordingly narrow (Fig. 7.2b).

Overall, we observe a smooth transition in the experimental PoP as a function of bath dimension from the wide exponential distribution for $D_B = 1$ to a narrower, Gaussian-like distribution for $D_B \gg 1$, in excellent agreement with the theoretical Erlang distribution over several orders-of-magnitude (Fig. 7.3c).

We do observe slight deviations from the Erlang distribution in Fig. 7.3c for $D_B \approx D$ as our derivation assumed $D \rightarrow \infty$. Still, we can calculate arbitrary moments of the PoP accounting for finite D , as we demonstrate explicitly for the second and third moments in the Section 7.4.5. In particular, we predict the variance of the PoP is

$$\sigma_{\text{rel}}^2 = \frac{1}{D_B} - \frac{1}{D}, \quad (7.2)$$

agreeing with the Erlang distribution variance for $D \rightarrow \infty$.

Importantly, the predictions of Eqs. 7.1 and 7.2 are not unique to the Rydberg Hamiltonian, but are instead universal for *any* quantum many-body system which exhibits HSE. To observe this, we plot $\sigma_{\text{rel}}^2 + 1/D$ for various systems expected to exhibit HSE including our experimental Rydberg system, and also numerical simulations of random unitary circuits (RUC), mixed field Ising models, and itinerant models like Bose-Hubbard systems (Fig. 7.4). In all cases, we find good agreement with the analytic prediction of Eq. (7.2). Later, we will introduce a rescaling of the bath dimension to treat systems at finite temperature or with conservation laws, with experimental and numerical support (see Fig. 7.8).

On the contrary, for failure cases which do not exhibit HSE (because they do not fulfill the no-resonance conditions), such as the transverse field Ising model [346] and systems of non-interacting spins, the prediction of Eq. (7.2) is violated (Fig. 7.4 inset). Specifically, the estimated scaling coefficient with bath dimension $-\langle D_B(\sigma_{\text{rel}}^2 + 1/D) \rangle$ where $\langle \cdot \rangle$ is an average over choices of D_B – is uniformly 1 for systems exhibiting HSE, but shows non-uniform and progressively greater deviations from 1 as a function of system size for non-ergodic cases.

In summary, we identify universality in the relative fluctuations of quantum systems interacting with an intrinsic bath under the condition of HSE, and we show the PoP for such systems is a universal Erlang distribution. When the bath is small, large quantum fluctuations dominate, but when the bath is large, quantum fluctuations are damped out, leaving only small, Gaussian fluctuations. We identify this smooth quantum-to-classical crossover as integral to the transition from quantum mechanical to statistical mechanical descriptions of generic, closed, quantum many-body systems as the number of bath configurations increases.

7.2 OPEN SYSTEM DYNAMICS

So far, we studied the case of closed system dynamics, where A reaches a thermal equilibrium due to interaction with its complement subsystem, B . Now we consider the open, composite system AB , and explore if similar universal predictions hold for noisy interactions with the external environment, E (Fig. 7.1a). Here we will study the PoP aggregated over all bitstrings at a fixed time instead of over measurement times; HSE entails these two pictures are equivalent [30, 71], as we verify numerically in the Fig. 7.9. Explicitly we focus on the global PoP distribution, not the local PoP measured in a subsystem.

To approach this task, we revisit our understanding of the Erlang distribution for closed system thermalization. Under HSE, the global bitstring probability, $\tilde{p}(z)$, is randomly sampled from an exponential distribution. The marginal probability, $\tilde{p}(z_A)$, is then a sum of D_B uncorrelated exponential random variables, each with equal weight, $\omega=1/D_B$, as the infinite temperature bath has no preferred configuration. This process leads precisely to the Erlang distribution, i.e. the Erlang distribution is the convolution of D_B exponential distributions [291] (Fig. 7.5a).

We can make a similar assumption for open system dynamics – that for each bath configuration the system samples an independent exponential distribution – but now allow for non-equal weights that are determined by the microscopic details of the extrinsic coupling mechanisms. As an example, we consider the case of RUC evolution where a single error can occur on a random qubit sometime during the evolution. For a circuit with N qubits run for a depth D , there are $K = N \times D$ potential error locations, and the circuit executes with a finite global fidelity of F . Thus, with a probability of F the circuit will execute without error and will sample bitstring probabilities from an exponential distribution. However, if an error does occur at one of the K locations, assuming the circuit is sufficiently scrambling [275] it will sample a separate *independent* exponential distribution. Without the ability to explicitly keep track of all errors, the aggregate distribution describes a sum of independent exponential variables with weights $\omega_1 = F$ and $\{\omega_2, \dots, \omega_{K+1}\} = (1-F)/K$ (Fig. 7.5b), which are related to the system's Kraus operators (Section 7.7.2).

When weights are allowed to be non-equal, we identify the result as a *hypoexponential distribution* [291, 347], or generalized Erlang distribution. Importantly, largely arbitrary noise channels coupling the quantum system to the external environment can be captured simply by adjusting the weight vector (Fig. 7.5b). To demonstrate this, we perform numerical simulations of RUC evolution and mixed field Ising model Hamiltonian evolution (Section 7.3). In both cases we apply either: 1) global depolarizing error, 2) local amplitude damping

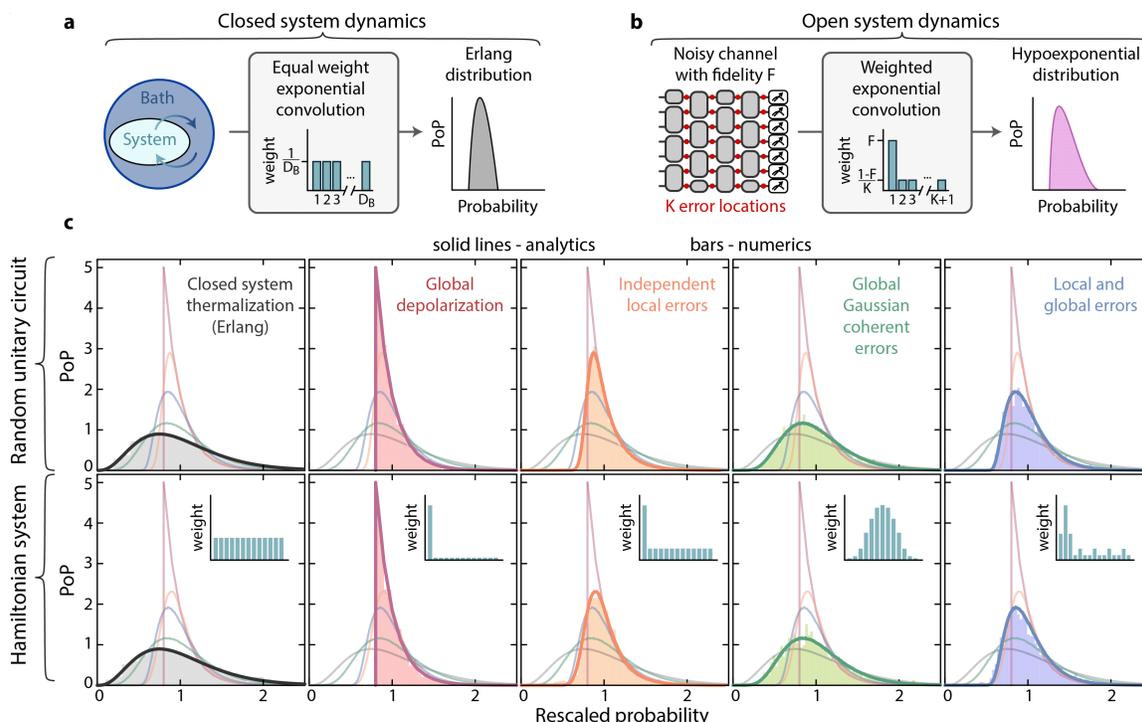


Figure 7.5: **Distinguishing noise channels using PoP distributions.** **a.** The Erlang distribution observed in the case of closed, infinite temperature thermalization can be alternatively derived as a convolution of independent exponential distributions with equal weights, one for each bath configuration. **b.** We extend this heuristic derivation to the case of open system dynamics, where the weights in the convolution are determined by the specifics of the noise channel, for example here a local error channel (see text), resulting in the PoP being a *hypoexponential distribution*. **c.** Hypoexponential distributions for four different noise channels each with distinct hypoexponential weights (insets, see Section 7.3 for model details), with noise strengths set to result in the same fidelity $F \approx 0.2$. Also shown is the Erlang distribution from closed system thermalization with an intrinsic bath of two qubits. PoP obtained from numerical simulations of RUC (top) and Hamiltonian evolution under a mixed field Ising model (bottom) show excellent agreement with their corresponding analytical predictions (thick, dark lines), but can easily be distinguished from predictions of other bath types (thin, faint lines), both intrinsic and extrinsic.

errors, 3) global Gaussian coherent errors, or 4) a combination of errors (2) and (3). For each, the error strength is tuned so the final fidelity is $F \approx 0.2$. By *coherent error*, we mean a static, shot-to-shot under/over rotation, which here we take to be Gaussian-distributed and applied to all qubits globally. We note that the weights for the global depolarizing channel are the same as those of the local error channel with $K \rightarrow \infty$, consistent with Ref. [348] (Fig. 7.14). Note, however, that the PoP for global coherent errors does not converge to the PoP of global depolarizing noise for increasing system sizes at fixed fidelity. Coherent errors could arise from non-Markovian Hamiltonian parameter variations [31, 349], or systematic

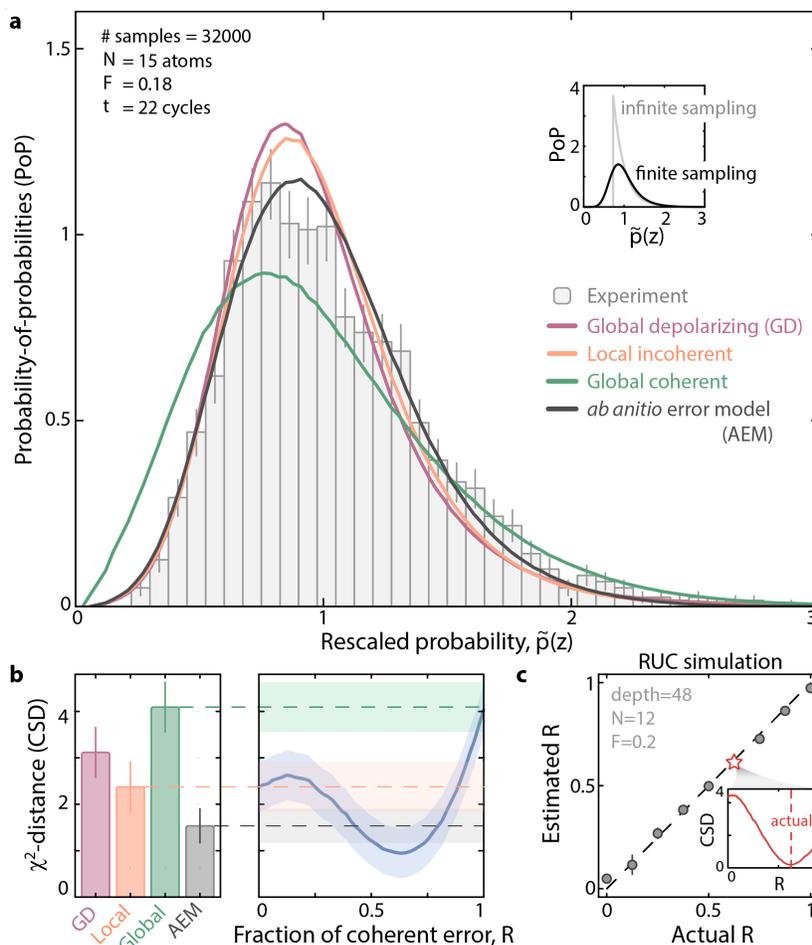


Figure 7.6: **Characterizing experimental noise.** **a.** We compare the PoP distribution obtained from the experiment following long-time evolution against various noise models, including our own *ab initio* error model. Model PoPs are nontrivially modified by finite sampling effects (inset, see Section 7.6). **b.** We compute the symmetric χ^2 -distance (CSD) between the experimental distribution and the model PoPs (left), including to a mixture error model composed of global coherent error and local incoherent error with varying global coherent error fraction, R . **c.** We perform noisy RUC simulations with varying R , and find that estimating R from the CSD minimum is well-correlated with the actual R .

gate miscalibrations [162, 350]; discriminating such noise sources from Markovian errors is crucial as they lead to different functional forms of fidelity decay which may be otherwise difficult to distinguish [31, 351], and their mitigation necessitates different models of error correction [227, 342–344].

In all cases, we see excellent agreement between the numerical PoP distributions and their corresponding analytical hypoexponential predictions, while being visually distinct from those of other noise channels, and from thermalization with an infinite temperature bath with $D_B = 4$ (Fig. 7.5c). Thus, on a qualitative level we can discriminate between a wide

range of noise models through many-body bitstring probability measurements, even when those models lead to an equal final fidelity.

We then turn to apply this method quantitatively to our experimental quantum simulator. We program the same time-independent dynamics as we studied in the case of intrinsic thermalization, but now evolve out to late times with $N=15$ atoms such that the global fidelity is $F \approx 0.18$, as estimated via the F_d cross-entropy variant [30, 31]. We then compare the experimental PoP distribution against the predictions for various noise models including our in-house *ab initio* noise model capturing all known physical noise sources [31]. The use of a finite number of samples modifies the predicted PoP distributions (inset of Fig. 7.6a, Section 7.6).

A qualitative evaluation (Fig. 7.6a) already suggests our *ab initio* noise model is the best match to the experiment. To quantify this statement, we compute the normalized symmetric χ^2 -distance (CSD) between the experiment and the model PoPs (Section 7.6, Fig. 7.6b), and find the distance is indeed minimized for the *ab initio* noise model.

Further, we apply this technique to error models with a mixture of global Gaussian coherent error and local incoherent error. We vary the fractional contribution of global coherent error, R , and analytically construct the model PoP to compare against (Section 7.5). We find the lowest CSD is found for $R \approx 2/3$. This technique is applicable to both analog quantum simulators and digital quantum processors. For instance, we simulate noisy RUC dynamics for different choices of R , and for each compare the R estimated from the CSD minimum versus the actual R (Fig. 7.6c). We find the estimated and actual values are well-correlated (which improves with deeper circuits, see Fig. 7.15) across the entire range of R .

Importantly, our analysis requires no circuit or Hamiltonian simulation, and makes no assumption about the underlying noise model besides its weight vector. However, this also means it cannot discriminate between noise models with the same weights, such as distinguishing local bit flips from local phase flips. Thus, it should be considered a tool for efficiently identifying experimentally-consistent models with bespoke noise distributions, which can then be further refined via additional measurements to distinguish different models with the same weights.

7.3 DETAILS OF EXPERIMENTS AND SIMULATIONS

We summarize details of the experiment and of the simulations which were used throughout this chapter.

7.3.1 Experimental details

For the $N = 10$ data at approximately infinite temperature shown in Figs. 7.2, 7.3 and 7.4, the Hamiltonian parameters are chosen to be: $\Omega/2\pi = 5.3$ MHz, $\Delta/2\pi = 0.6$ MHz, $C_6/2\pi = 254$ GHz $\times\mu\text{m}^6$, and $a = 3.77$ μm (the lattice spacing is calibrated with a laser-based ruler [28]). For the $N = 15$ data at approximately infinite temperature shown in Fig. 7.6, the Hamiltonian parameters are the same, but with $\Omega/2\pi = 5.4$ MHz. For the $N = 10$ data at finite temperature in Fig. 7.8, the parameters are the same, but with $\Omega/2\pi = 5.3$ MHz, and $\Delta/2\pi = 4.6$ MHz.

Through measurement, we accrue a set of time-labeled bitstrings, $\{z\}$. When normalizing bitstring probabilities, we obtain the temporal long time average directly from the *diagonal ensemble* [30, 305, 326],

$$\hat{\rho}_d = \sum_n |c_n|^2 |E_n\rangle\langle E_n|, \quad (7.3)$$

where $|E_n\rangle$ are energy eigenstates, and $c_n = |\langle E_n|\Psi_0\rangle|^2$ are overlaps with the initial state. Any observable, including bitstring probabilities, of $\hat{\rho}_d$ is equal to the same observable obtained from an infinite time average:

$$p_{\text{avg}}(z) = \langle z|\hat{\rho}_d|z\rangle = \int_0^\infty dt |\langle z|\Psi(t)\rangle|^2. \quad (7.4)$$

A major feature of our system Hamiltonian is the Rydberg blockade mechanism [36], whereby the energy eigenspectrum is naturally stratified into a number of widely spaced sectors. In particular, the nearest-neighbor interaction is $\approx 20\times$ greater than the next largest energy scale, so cases where neighboring pairs of atoms are both excited to the Rydberg state are greatly suppressed. The resulting effective Hilbert-space dimension, D , of a system with N qubits is given by $\text{Fib}(N + 2) \approx 1.6^N$, where Fib is the Fibonacci function. To be clear, statements in this work that the Rydberg system thermalizes at infinite effective temperature refer only to thermalization within the blockade-satisfying subspace, not across the entire Hilbert space.

For $N = 10$, we find $\approx 99\%$ of all experimental bitstrings are in the blockade-satisfying energy sector at short times; we typically refine $\{z\}$ by discarding all realizations not in this sector. For the late-time $N = 15$ data the number of blockade-satisfying bitstrings drops to $\approx 75\%$, which we attribute mainly to spontaneous decay [26], though leave careful investigation to future work.

The effective dimension of the intrinsic bath, D_B , is altered by the blockade constraint. Specifically, the blockade must be accounted for at the boundary between subsystems

A and B , such that D_B then depends on the bitstring measured in A . As an example, consider the system A as the leftmost atoms in a one dimensional array with ten total atoms, and B as the remaining $N_B = 9$ atoms, such that the global bitstring is written as $z = z_A z_B$. For the subsystem bitstring $z_A = 0$, there are no possible blockade violation across the boundary between A and B , and so the bath dimension is given by $D_B(z_A = 0) = \text{Fib}(N_B + 2) = 89$. However, for the subsystem bitstring $z_A = 1$, the blockade constraint forbids the leftmost qubit in B from also being a 1. Thus the bath dimension is $D_B(z_A = 1) = \text{Fib}((N_B - 1) + 2) = 55$. In general, we allow A to be any contiguous block of atoms in the array (meaning B may be discontinuous), and for instance in the analysis of Fig. 7.3 we aggregate over bitstrings in A with equal effective bath dimension.

Error bars on the experimental data in Fig. 7.4 are 1σ standard error of the variance taken over the different choices of z_A at a given D_B . The variance for the random unitary circuit in Fig. 7.4B is taken as the median over 1000 circuit realizations; the actual circuit being simulated is composed of SU(4) gates in a brickwork pattern. In Fig. 7.6, The error model prediction is calculated via a Monte Carlo wavefunction approach [29, 31, 114] from which bitstrings are sampled at an equal number as the experiment. Experimental error bars in Fig. 7.6 are analytic based on the Poisson sampling errors: for a histogram bin with m counts, its standard error is \sqrt{m} (Section 7.6.3).

7.3.2 Simulation details

Here we summarize the parameters used in numerical simulations throughout. In general, PoP distributions from numerics are aggregated over several hundred cycles of evolution, except where otherwise noted.

Error free simulations

Rydberg: Rydberg simulation parameters are the same as for experiment, as given in a previous section.

MFIM: For the mixed field Ising model simulations in Fig. 7.4 and Fig. 7.8, we utilize the following Hamiltonian [322]:

$$\hat{H}/\hbar = \sum_{i=1}^N (h_x \hat{S}_i^x + h_y \hat{S}_i^y - \zeta \hat{S}_i^z) + \sum_{i=1}^N J \hat{S}_i^x \hat{S}_{i+1}^x, \quad (7.5)$$

where $h_x = 0.41$, $h_y = 0.45$, and $J = 1$. The field $\zeta = 0$ corresponds to infinite temperature evolution and $\zeta = 0.86$ to finite temperature evolution. In Fig. 7.13, we instead set $\zeta = 0$ but tune the initial state as $|\Psi_0\rangle = \left[\exp(i\theta \hat{S}_i^x) |0\rangle \right]^{\otimes N}$ for various angles θ .

TFIM: For the transverse field Ising model simulations in the inset of Fig. 7.4 we consider a one-dimensional open chain under the following Hamiltonian:

$$\hat{H}/\hbar = - \sum_i \hat{S}_i^x - 3 \sum_{i \neq j} \hat{S}_i^z \hat{S}_j^z. \quad (7.6)$$

We use as an initial state of $|\Psi(t=0)\rangle = |+\rangle^{\otimes N}$.

Non-interacting spins: Non-interacting spins studied in the inset of Fig. 7.4 are initialized in their ground state, and then evolved with equal Rabi frequency and no detuning, thus undergoing disconnected, correlated Rabi oscillations.

Bose Hubbard: We simulate the Bose-Hubbard model with N particles on a one-dimensional open chain with length M :

$$\hat{H}/\hbar = -\Omega \sum_{j=1}^{M-1} \left(\hat{b}_j^\dagger \hat{b}_{j+1} + \text{h.c.} \right) + \frac{U}{2} \sum_{j=1}^M \hat{n}_j (\hat{n}_j - 1), \quad (7.7)$$

where $\hat{n}_j = \hat{b}_j^\dagger \hat{b}_j$. We choose parameter values $(\Omega, U) = (1, 2.87)$. In our simulations, we have $M = N$, and we take the initial state to be the one with one particle at each site, i.e. $|1, 1, \dots, 1\rangle$.

Open system simulations

Rydberg Hamiltonian dynamics: Our *ab initio* error model for Rydberg Hamiltonian dynamics has been described previously [29], and shown to be accurate to experiment in several regimes [26, 29, 31]. We note that many of these noise sources generically are highly non-Markovian with important impact on the many-body fidelity [31]. Further, spontaneous decay channels are potentially non-unital, and predominantly dominated by leakage errors out of the qubit subspace [26].

Spin chain Hamiltonian dynamics: *Amplitude damping* – To obtain the bottom row of Fig. 7.5, we simulate the open system dynamics of the mixed Field Ising model [Eq. (7.5)] in the presence of amplitude damping, or spontaneous decay $|0\rangle\langle 1|$ errors, with the quantum trajectory method [114, 352]. In all trajectory methods, we perform a sufficient number of trajectories, typically 1000 for a system size of $N = 12$, to see convergence of the PoP.

Global coherent errors – To simulate global coherent errors, we again perform a trajectory method, where in each trajectory, we sample the parameter J in Eq. (7.5) from a normal distribution $N(J_0, \sigma_J^2)$, with $J_0 = 1$ and σ_J tuned to achieve the desired fidelity at a given time.

Mixed error sources – To simulate both amplitude damping and global coherent errors, we simply use a quantum trajectory method with both noise sources enabled.

Random unitary circuits: *Local errors* – We simulate the open system dynamics of a random unitary circuit (RUC) in the presence of amplitude damping errors. We evolve the density matrix by applying a layer of random unitary gates, before applying Kraus operators of the amplitude damping channel (for the top row of Fig. 7.5) independently at every qubit.

Global coherent errors – In order to model global coherent errors in a RUC, we parameterize each ideal random two-qubit gate as an $SU(4)$ rotation about some axis and with some angle. In each run we model the angles of each gate to be systematically over- or under-rotated by a factor which is constant over all gates within each run, but sampled from a normal distribution over many runs. For the mixed error sources in Fig. 7.5, we perform this stochastic over- and under-rotation with the density matrix evolution method outlined above.

Mixed error sources – For the mixed error sources in Fig. 7.6, we use a quantum trajectory method with both bitflip errors and stochastic over- and under-rotation errors (outlined above).

7.4 DERIVING THE ERLANG DISTRIBUTION

In this work, we studied closed system dynamics and predicted the distribution functions of bitstring probabilities (or, probability-of-probabilities) in a subsystem which is strongly coupled to the rest of the system that serves as a bath. Our main quantitative prediction was Eq. (7.1), the Erlang distribution, for such systems at infinite effective temperature and large total system sizes. Subsequently, we predicted the variance of the probability-of-probabilities, including $1/D$ corrections from the Erlang value.

Here, we provide a rigorous derivation of the Erlang distribution. First we formally define the probability-of-probabilities (PoP) distribution, then describe how Hilbert-space ergodicity entails that the PoP for global, isolated quantum systems is the Porter-Thomas distribution, with more formal details left to our recent theoretical work Ref. [71]. Then, we use this fact to derive that the Erlang distribution describes the PoP of subsystems thermalizing at infinite effective temperature, during closed system dynamics and in the large system size limit. In more technical sections, we then relax some of our assumptions and show how arbitrary order moments of the PoP may be calculated even at finite effective temperatures and finite system sizes. We also note that we provide a different, more heuristic, derivation of the Erlang distribution in Section 7.5.3.

7.4.1 The probability-of-probabilities distribution

The PoP distribution for a single bitstring, z , is a histogram of the bitstring probabilities, $p(z, t)$, formally over infinite time. For a given probability, p_0 , we write it as (where we drop the z for notational simplicity)

$$P(p_0) = \mathbb{E}_t[\delta(p(t) - p_0)] \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \delta(p(t) - p_0) dt. \quad (7.8)$$

The PoP encodes all temporal fluctuations of the time-series $p(t)$ as its moments,

$$\int_0^1 p_0^k P(p_0) dp_0 = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \int_0^1 p_0^k \delta(p(t) - p_0) dp_0 dt \quad (7.9)$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T p(t)^k dt \equiv \mathbb{E}_t[p(t)^k]. \quad (7.10)$$

The PoP is defined here for a single bitstring over all evolution times, but it can alternatively be defined for a single time over all bitstrings, a distinction which we will revisit below.

In this study, we are chiefly interested in studying *relative* temporal fluctuations, i.e. normalizing the time-series $p(t)$ to have mean 1. For random-circuit dynamics, this is simply the rescaling $\tilde{p} \equiv Dp$, while in Hamiltonian dynamics, this is $\tilde{p}(z) \equiv p(z)/p_{\text{avg}}(z)$, where $p_{\text{avg}}(z) \equiv \mathbb{E}_t[p(z, t)]$ is the time-averaged probability of z . This accordingly rescales the PoPs, which we will write as $P(\tilde{p})d\tilde{p}$ (occasionally omitting $d\tilde{p}$). The $p_{\text{avg}}(z)$ may be found from the diagonal ensemble:

$$\hat{\rho}_d = \sum_n |c_n|^2 |E_n\rangle\langle E_n|, \quad (7.11)$$

where $|E_n\rangle$ are energy eigenstates, and $c_n = |\langle E_n | \Psi_0 \rangle|^2$ are overlaps with the initial state. Any observable, including bitstring probabilities, of $\hat{\rho}_d$ is equal to the same observable obtained from an infinite time average:

$$p_{\text{avg}}(z) = \langle z | \hat{\rho}_d | z \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt |\langle z | \Psi(t) \rangle|^2. \quad (7.12)$$

7.4.2 The Porter-Thomas distribution from Hilbert-space ergodicity

The chief object of our interest is the *temporal ensemble*, the set of pure states, $|\Psi(t)\rangle$, evaluated at different times t for evolution under a time-independent Hamiltonian, \hat{H} . Explicitly, if we discretize the time evolution with infinitesimal step size, δ , and define the time $t_j = \delta j$ for $j \in \mathbb{Z}_+$, the temporal ensemble is the set $\{|\Psi(t_j)\rangle\}_{j \in \mathbb{Z}_+}$. The initial state sets a well-defined energy, which then constrains the temporal ensemble to a

subset of Hilbert space obeying energy conservation. Specifically, the Schrödinger equation conserves the population $|c_n|^2$ of each energy eigenstate $|E_n\rangle$, where $|c_n|^2 \equiv |\langle \Psi_0 | E_n \rangle|^2$. This imposes a total of D (the Hilbert space dimension) constraints, and the only remaining degrees of freedom are the complex phases $\arg(\langle E_n | \Psi(t) \rangle)$.

The temporal ensemble has been previously studied [30, 71, 345], where at late times it was connected to the *random phase ensemble*, the ensemble of quantum states with well-defined amplitudes in the energy eigenbasis, but random phases. This connection depends on so-called *no-resonance conditions* [336], which are colloquially the absence of high-order degeneracies in the energy eigenspectrum. More concretely, a Hamiltonian \hat{H} satisfies the k -th no-resonance condition if for any two sets of k eigenvalues $\{E_{\alpha_j}\}_{j=1}^k$ and $\{E_{\beta_j}\}_{j=1}^k$ the equation

$$E_{\alpha_1} + E_{\alpha_2} + \cdots + E_{\alpha_k} = E_{\beta_1} + E_{\beta_2} + \cdots + E_{\beta_k} \quad (7.13)$$

is true if and only if the sets of indices $(\alpha_1, \alpha_2, \dots, \alpha_k)$ and $(\beta_1, \beta_2, \dots, \beta_k)$ are equal up to reordering. For example, the 2nd-no resonance condition states that the relation $E_a + E_b = E_c + E_d$ is true if and only if $a = c, b = d$ or $a = d, b = c$. That is, there are no “accidental resonances.” Existence of k -th no-resonance conditions for any order k guarantee the temporal ensemble is equivalent to the random phase ensemble in the infinite time limit [71]; as shown in Ref. [71] this then is equivalent to the statement that Hilbert-space ergodicity holds.

We are interested in the k -th moment of the temporal ensemble, which, up to corrections which are exponentially small in the system size, is equal to [71]

$$\mathbb{E}_t[|\Psi(t)\rangle\langle\Psi(t)|^{\otimes k}] \approx \hat{\rho}_d^{\otimes k} \sum_{\sigma \in S_k} \text{Perm}(\sigma), \quad (7.14)$$

where $\hat{\rho}_d$ is the diagonal ensemble [Eq. (7.11)] and $\text{Perm}(\sigma)$ is the permutation operator acting on k copies of the Hilbert space, which, for a permutation σ in the symmetric group S_k , takes any configuration $|i_1, i_2, \dots, i_k\rangle \in \mathcal{H}^{\otimes k}$ and maps it to $|i_{\sigma(1)}, i_{\sigma(2)}, \dots, i_{\sigma(k)}\rangle$. Explicitly, the permutation operator has $k!$ elements. As an example, for $k = 2$

$$\mathbb{E}_t[|\Psi(t)\rangle\langle\Psi(t)|^{\otimes 2}] \approx \hat{\rho}_d^{\otimes 2} (\hat{I} + \hat{S}), \quad (7.15)$$

where \hat{I} is the identity, \hat{S} is the swap operator between two copies of Hilbert space, and $\hat{\rho}_d$ is the diagonal ensemble (Eq. (7.11)). Importantly, this is the same as the k -th moment of the Haar ensemble [285], taking $\hat{\rho}_d = \hat{I}/D$.

To study fluctuations of bitstring probabilities, we simply contract Eq. (7.14) with $\langle z |^{\otimes k}$ and $|z\rangle^{\otimes k}$, where each of the $k!$ permutations have equal contribution $p_{\text{avg}}(z)^k$. This contraction

is equivalent to the k -th moment of the probability time series [71]

$$\mathbb{E}_t[p(z, t)^k] \approx k! p_{\text{avg}}(z)^k. \quad (7.16)$$

These are the same as the moments of the Porter-Thomas (exponential) distribution with mean value $p_{\text{avg}}(z)$. Thus, upon rescaling $\tilde{p}(z, t) = p(z, t)/p_{\text{avg}}(z)$, the resulting distribution is exactly the Porter-Thomas with PoP

$$P(\tilde{p})d\tilde{p} = \exp(-\tilde{p})d\tilde{p}. \quad (7.17)$$

7.4.3 The Erlang distribution from Hilbert-space ergodicity

From Eq. (7.14), we can also derive that the rescaled marginal distributions $p(z_A)/p_{\text{avg}}(z_A)$ follow the Erlang distribution.

Proof. We consider the k -th cumulants of $p(z_A)$, which is related to the k -th moment, but subtracts off trivial contributions from lower order moments. The k -th cumulant of the Erlang distribution can be simply calculated as $\kappa_k = (k-1)!/D_B^{k-1}$, since the Erlang distribution is the convolution of multiple exponential distributions [291], and cumulants of convolutions are additive [353].

To calculate the k -th cumulant we consider the permutations in Eq. (7.14). The permutation operator swaps between different copies of the Hilbert space, and can be represented by a directed graph indicating the action of the permutation on k elements. Importantly, this graph can be disconnected, e.g. for three copies of the Hilbert space the permutation operator can leave one unchanged while swapping the other two. However, in this case the permutation will only contain terms which are related to lower-order moments. The cumulant explicitly subtracts off such lower order moments, and thus it is related only to *cyclic* permutations, i.e. permutations with only a single cycle [353].

Applying this to our observable of interest, we write the general formula for the k -th cumulant of the time series of $p(z_A)$, κ_k , as:

$$\kappa_k = \sum_{\sigma \in C_k} \sum_{z_B^{(1)}} \cdots \sum_{z_B^{(k)}} \left(\bigotimes_{i=1}^k \langle z_A, z_B^{(i)} | \right) \hat{\rho}_d^{\otimes k} \text{Perm}(\sigma) \left(\bigotimes_{i=1}^k | z_A, z_B^{(i)} \rangle \right) \quad (7.18)$$

where \bigotimes represents a repeated Kronecker product, and C_k is the subset of S_k containing only cyclic permutations, which importantly has $(k-1)!$ elements. This is a general statement for cumulants of a local bitstring probability, $p(z_A)$, in the

large total system size limit; it is much like Eq. (7.16), but is subselected only to the cyclic permutations and with the inclusion of sums over different choices of z_B in the contraction.

To simplify this equation, it is helpful to first consider its lowest non-trivial form for $k = 2$, for which the only contributing permutation is the swap operator, \hat{S} , with explicit steps

$$\kappa_2 = \sum_{z_B, z'_B} (\langle z_A, z_B | \otimes \langle z_A, z'_B |) \hat{\rho}_d^{\otimes 2} \hat{S} (|z_A, z_B\rangle \otimes |z_A, z'_B\rangle) \quad (7.19)$$

$$= \sum_{z_B, z'_B} \langle z_A, z_B | \hat{\rho}_d | z_A, z'_B \rangle \langle z_A, z'_B | \hat{\rho}_d | z_A, z_B \rangle \quad (7.20)$$

$$= \sum_{z_B} \langle z_A, z_B | \hat{\rho}_d | z_A, z_B \rangle^2 + \sum_{z_B \neq z'_B} |\langle z_A, z_B | \hat{\rho}_d | z_A, z'_B \rangle|^2 \quad (7.21)$$

$$\approx \sum_{z_B} \langle z_A, z_B | \hat{\rho}_d | z_A, z_B \rangle^2. \quad (7.22)$$

In the third line, we have split the sum over different choices of z_B and z'_B into cases where $z_B = z'_B$ and cases where $z_B \neq z'_B$. In the final line, we have made the infinite temperature approximation that $\hat{\rho}_d \approx \hat{I}/D$ – up to small variations along the diagonal due to the varied $p_{\text{avg}}(z)$ – which nullifies the second term, where $z_B \neq z'_B$. In other words, we assume off-diagonal terms are negligible.

Carrying out a similar calculation – and again only keeping terms where all z_B are equal – for the k -th cumulant yields

$$\kappa_k \approx (k-1)! \sum_{z_B} \langle z_A, z_B | \hat{\rho}_d | z_A, z_B \rangle^k, \quad (7.23)$$

where the coefficient $(k-1)!$ comes precisely from the number of elements in the set of cyclic permutations. To be explicit, by selecting cases where all z_B are equal (because of the infinite temperature condition) the cyclicity of the permutations does not actually matter for our proof, beyond the fact that there are $(k-1)!$ of them.

Then to further simplify Eq. (7.23), we assume that any z_B dependence can be ignored and we can approximate

$$\sum_{z_B} \langle z_A, z_B | \hat{\rho}_d | z_A, z_B \rangle^k \approx \langle z_A | \text{tr}_B(\hat{\rho}_d) | z_A \rangle^k / D^{k-1} \quad (7.24)$$

$$= p_{\text{avg}}(z_A)^k / D^{k-1}. \quad (7.25)$$

This is an assumption that the basis $\{z_B\}$ is *energy non-revealing*, see Section 7.5.3.

Thus, when considering cumulants of relative quantities, i.e dividing the k -th cumulant by $p_{\text{avg}}(z_A)^k$, we find the k -th cumulant is $(k-1)!/D_B^{k-1}$. This is precisely equal to the k -th cumulant of the Erlang distribution. Since all cumulants match, all moments must match too, and we conclude that the marginal probabilities $p(z_A)$ follow an Erlang distribution for systems thermalizing at infinite temperature, up to corrections exponentially small in the system size.

In this derivation, we have made two approximations. First, we have assumed that the basis z_B is energy non-revealing, and second, we have assumed that the off-diagonal terms of the form $\langle z_A z_B | \hat{\rho}_d | z_A z'_B \rangle$ are negligible. Remarkably, we may perform an exact calculation of the k -th moments $\mathbb{E}_t[p(z_A, t)^k]$ to estimate the effect of these approximations. We outline these more technical results below. Among other results, this allows us to obtain the $1/D$ correction in the expression of the relative variance σ_{rel}^2 .

7.4.4 Second Hamiltonian twirling identity

Our primary tool will be the “second Hamiltonian twirling identity,” which describes “twirling” under a family of channels $U_t \equiv \exp(-iHt)$ formed by evolution under a fixed Hamiltonian H , over different times t , and was established in Refs. [30, 71]. This is analogous to a conventional twirling channel composed of Haar random unitaries and its results bear many similarities [354].

Assuming the second no-resonance condition, the twirling channel, τ_2 , acting on any two-copy operator $\hat{A} : \mathcal{H}^{\otimes 2} \rightarrow \mathcal{H}^{\otimes 2}$ is given as

$$\tau_2 = \mathbb{E}_t \left[\hat{U}_t^{\otimes 2} \hat{A} \hat{U}_{-t}^{\otimes 2} \right] \quad (7.26)$$

$$\equiv \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt \hat{U}_t^{\otimes 2} \hat{A} \hat{U}_{-t}^{\otimes 2} \quad (7.27)$$

$$= \sum_{i,j,k,l} \left[\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt e^{-i(E_i + E_j - E_k - E_l)t} \right] |E_i, E_j\rangle \langle E_i, E_j| \hat{A} |E_k, E_l\rangle \langle E_k, E_l| \quad (7.28)$$

$$= \sum_{i,j,k,l} \delta(E_i + E_j - E_k - E_l) |E_i, E_j\rangle \langle E_i, E_j| \hat{A} |E_k, E_l\rangle \langle E_k, E_l| \quad (7.29)$$

$$= \sum_{E, E'} \left[|E, E'\rangle \langle E, E'| \hat{A} |E, E'\rangle \langle E, E'| + |E, E'\rangle \langle E, E'| \hat{A} |E', E\rangle \langle E', E| \right] \quad (7.30)$$

$$- \sum_E |E, E\rangle \langle E, E| \hat{A} |E, E\rangle \langle E, E|.$$

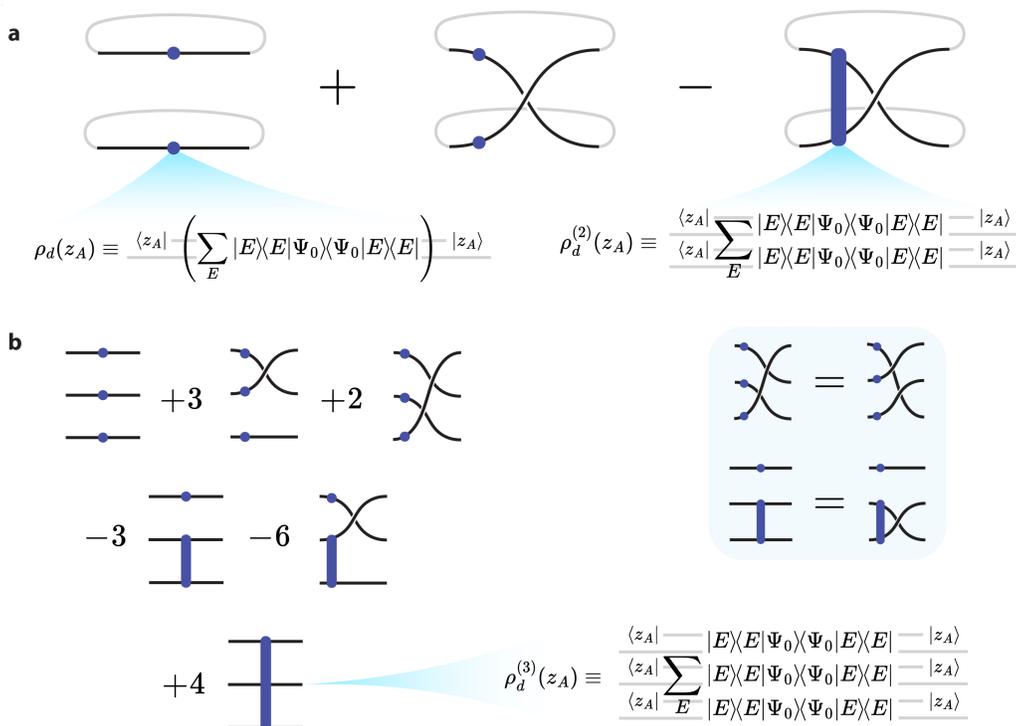


Figure 7.7: Diagrams for subsystem bitstring probability fluctuations. We diagrammatically organize our computations for the second and third moments $\mathbb{E}_t[p(z_A, t)^k]$, which we show are equal to the low moments of the Erlang distribution at infinite temperature. **a.** We organize the terms in the second moment of the temporal ensemble diagrammatically, which translates to the equation $\mathbb{E}_t[p(z_A, t)^2] = \text{tr}[\hat{\rho}_d(z_A)]^2 + \text{tr}[\hat{\rho}_d(z_A)^2] - \text{tr}[\hat{\rho}_d^{(2)}(z_A)]$ [Eq. (7.40)]. The blue dot represents $\hat{\rho}_d(z_A)$, which is an operator on \mathcal{H}_B , and can be interpreted as the initial state, dephased in the energy eigenbasis, and projected onto a subsystem bitstring $|z_A\rangle$. The blue bar represents $\hat{\rho}_d^{(2)}(z_A)$, which acts on the double Hilbert space $\mathcal{H}_B^{\otimes 2}$, and is a subleading correction term to the first two terms. Each diagram is ultimately traced out (gray lines), which we subsequently omit for simplicity. **b.** The diagrammatic framework greatly simplifies the computation of the third moment [Eq. (7.51)]. The six permutations of S_3 give rise to three terms, since permutations with the same cycle type produce the same terms (blue box). Correction terms may be expressed in terms of diagrams with blue bars. Furthermore, diagrams where the lines with a blue bar are swapped are equivalent. Therefore, there are only two inequivalent diagrams involving $\hat{\rho}_d^{(2)}(z_A)$. The number of equivalent diagrams, along with double-counting coefficients from the twirling identity gives rise to the coefficients -3 and -6 . Finally, there is a higher-order diagram that accounts for triple-counting, involving $\hat{\rho}_d^{(3)}(z_A)$.

In Eq. (7.30), we applied the no-resonance condition to simplify our sum over four indices to a sum over two indices. The last term comes from the case where $i = j = k = l$, which is double counted in the first sum in Eq. (7.30). As discussed in Ref. [30], the second no-resonance condition is violated whenever there are energy degeneracies, which may be protected by non-Abelian symmetries. However, when the operator \hat{A} is specialized to $\hat{A} = \hat{\rho}_0^{\otimes 2}$ (as we do below), the degenerate energy eigenspace is projected onto a single state, and we may ignore such degeneracies: our analysis is valid as long as the second no-resonance condition *modulo degeneracies* applies. To aid their application, diagrammatic methods have been developed to represent the twirling identities [30, 299] (Fig. 7.7).

We specialize our results to $\hat{A} = \hat{\rho}_0^{\otimes 2}$, where $\hat{\rho}_0 = |\Psi_0\rangle\langle\Psi_0|$ is the initial pure state. This object is the second moment of the temporal ensemble, which we will use to study the variance of observables over time. Eq. (7.30) becomes

$$\text{LHS} = \mathbb{E}_t \left[\hat{U}_t^{\otimes 2} \hat{\rho}_0^{\otimes 2} \hat{U}_{-t}^{\otimes 2} \right] = \mathbb{E}_t [|\Psi(t)\rangle\langle\Psi(t)|^{\otimes 2}] \quad (7.31)$$

$$\begin{aligned} \text{RHS} = & \left[\left(\sum_E |E\rangle\langle E| \hat{\rho}_0 |E\rangle\langle E| \right) \otimes \left(\sum_{E'} |E'\rangle\langle E'| \hat{\rho}_0 |E'\rangle\langle E'| \right) \right] (\hat{I} + \hat{S}) \\ & - \sum_E |E, E\rangle\langle E, E| \hat{\rho}_0^{\otimes 2} |E, E\rangle\langle E, E|. \end{aligned} \quad (7.32)$$

Here $\hat{S} = \sum_{E, E'} |E, E'\rangle\langle E', E|$ is the swap operator between the two copies of Hilbert space. Thus, we arrive at the following simplified expression:

$$\mathbb{E}_t [|\Psi(t)\rangle\langle\Psi(t)|^{\otimes 2}] = \hat{\rho}_d^{\otimes 2} (\hat{I} + \hat{S}) - \hat{\rho}_d^{(2)}, \quad (7.33)$$

where $\hat{\rho}_d = \sum_E |E\rangle\langle E| \hat{\rho}_0 |E\rangle\langle E|$ is the diagonal ensemble (Eq. (7.11)) and $\hat{\rho}_d^{(2)} = \sum_E |E, E\rangle\langle E, E| \hat{\rho}_0^{\otimes 2} |E, E\rangle\langle E, E|$ is a two-copy generalization. The inclusion of the $\hat{\rho}_d^{(2)}$ term is crucially different compared to Eq. (7.14).

7.4.5 Universal variance including finite-size corrections

We study the variance of bitstring probabilities by using explicit expressions for the second moment of the temporal ensemble when Hilbert-space ergodicity applies [30, 71]. Specifically, we study the relative temporal fluctuation $\sigma_{\text{rel}}^2(z_A)$, and we find an expression that has a natural interpretation as an effective dimension of the finite-temperature bath.

We define the relative temporal fluctuation $\sigma_{\text{rel}}^2(z_A)$ of the probability of a particular bitstring z_A in subsystem A as follows:

$$\sigma_{\text{rel}}^2(z_A) \equiv \frac{\mathbb{E}_t [p(z_A, t)^2] - \mathbb{E}_t [p(z_A, t)]^2}{\mathbb{E}_t [p(z_A, t)]^2} = \frac{\mathbb{E}_t [p(z_A, t)^2]}{\mathbb{E}_t [p(z_A, t)]^2} - 1. \quad (7.34)$$

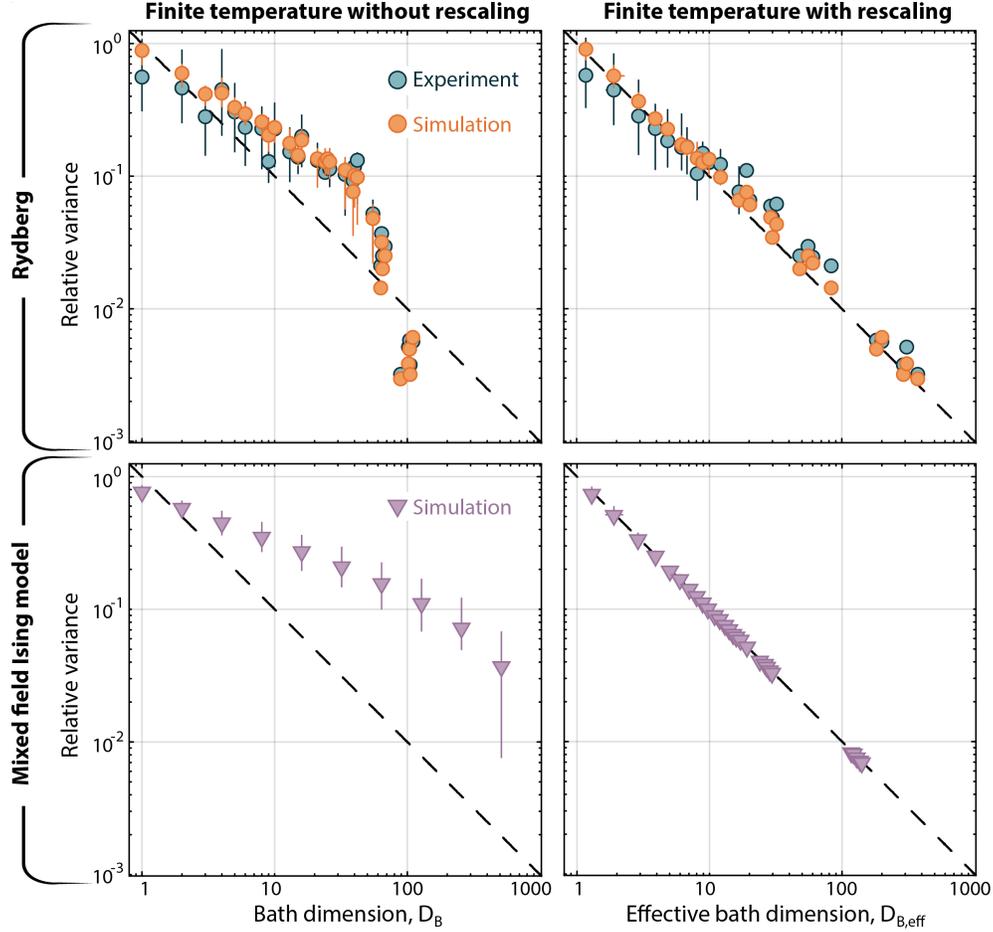


Figure 7.8: **Rescaling by effective bath dimension for finite temperature systems.** We show relative variance plots for Rydberg (both experiment and simulation) as well as for mixed field Ising model (simulation). In each case the quench evolution is at an finite effective temperature. In the left plots we show the relative variance (without the addition of a $1/D$ term) versus the bath dimensional, showing non-universal behavior. In the right plots we show the relative variance (plus the $1/D_\beta$ term, Eq. 7.44) versus the effective bath dimension, $D_{B,\text{eff}}$ (Eq. (7.43)), which recovers the scaling expected from Eq. (7.42). In all plots, the markers are averaged over all bitstring realizations with an effective bath dimension in an interval from $[i - 0.5, i + 0.5)$, for $i \in \mathbb{Z}$, and error bars are the standard deviations over these realizations.

Here $p(z_A, t) = \text{tr}[(|z_A\rangle\langle z_A| \otimes \hat{I}_B) |\Psi(t)\rangle\langle\Psi(t)|]$ is the marginal probability of measuring z_A at time t , where \hat{I}_B is the identity operator acting on subsystem B . The denominator of Eq. (7.34) can be simplified to

$$\mathbb{E}_t[p(z_A, t)] = \text{tr} \left[|z_A\rangle\langle z_A| \otimes \hat{I}_B \left(\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt |\Psi(t)\rangle\langle\Psi(t)| \right) \right] \quad (7.35)$$

$$= \langle z_A | \text{tr}_B(\hat{\rho}_d) | z_A \rangle \quad (7.36)$$

$$\equiv \text{tr}_B[\hat{\rho}_d(z_A)], \quad (7.37)$$

where in the second-last equality $\text{tr}_B(\hat{\rho}_d)$ is a mixed state on subsystem A , and in the last equality we instead consider the quantity $\hat{\rho}_d(z_A) \equiv (|z_A\rangle\langle z_A| \otimes \hat{I}_B)\hat{\rho}_d(|z_A\rangle\langle z_A| \otimes \hat{I}_B)$, which is a mixed state on the subsystem B .

Meanwhile, the numerator of Eq. (7.34) can be evaluated from the second moment of the temporal ensemble [30, 71]. This yields three terms: the permutation terms from Eq. (7.14), as well as a correction term illustrated in Fig. 7.7.

$$\mathbb{E}_t[p(z_A, t)^2] = \mathbb{E}_t \left[\left(\langle z_A | \text{tr}_B (|\Psi(t)\rangle\langle\Psi(t)|) |z_A\rangle \right)^2 \right] \quad (7.38)$$

$$= \text{tr} \left[(|z_A\rangle\langle z_A| \otimes \hat{I}_B)^{\otimes 2} \hat{\rho}_d^{\otimes 2} \right] + \text{tr} \left[(|z_A\rangle\langle z_A| \otimes \hat{I}_B)^{\otimes 2} \hat{\rho}_d^{\otimes 2} \hat{S} \right] \quad (7.39)$$

$$\begin{aligned} & - \text{tr} \left[(|z_A\rangle\langle z_A| \otimes \hat{I}_B)^{\otimes 2} \hat{\rho}_d^{(2)} \right] \\ & = \text{tr}_B[\hat{\rho}_d(z_A)]^2 + \text{tr}_B[\hat{\rho}_d(z_A)^2] - \sum_E q_{z_A}^2(E) \end{aligned} \quad (7.40)$$

The inclusion of $\hat{\rho}_d^{(2)}$ prevents double counting of certain terms in the sum of permutations (see Fig. 7.7), and is explicitly included here while previously (e.g. Eq. (7.14)) we disregarded it as negligible in the large system size limit.

Combining these results, we obtain the relative fluctuation as follows:

$$\sigma_{\text{rel}}^2(z_A) = \frac{\text{tr}_B[\hat{\rho}_d(z_A)]^2 + \text{tr}_B[\hat{\rho}_d(z_A)^2] - \sum_E q_{z_A}^2(E)}{\text{tr}_B[\hat{\rho}_d(z_A)]^2} - 1 \quad (7.41)$$

$$\equiv \frac{1}{D_{B,\text{eff}}} - \frac{1}{D_\beta}, \quad (7.42)$$

where

$$D_{B,\text{eff}}^{-1} \equiv \frac{\text{tr}_B[\hat{\rho}_d(z_A)^2]}{\text{tr}_B[\hat{\rho}_d(z_A)]^2} \quad (7.43)$$

$$D_\beta^{-1} \equiv \frac{\sum_E q_{z_A}^2(E)}{\text{tr}_B[\hat{\rho}_d(z_A)]^2}. \quad (7.44)$$

Here $D_{B,\text{eff}}^{-1}$ is the purity of the conditional reduced density matrix $\hat{\rho}_d(z_A)/\text{tr}_B[\hat{\rho}_d(z_A)]$. It measures the effective size of the bath, hence we dub it the *effective bath dimension*. Meanwhile, D_β^{-1} , is the purity of the classical distribution [30] (of size D) over the energy eigenstates, which has typical value $O(1/D)$. At infinite temperature, we can approximate

$\hat{\rho}_d(z_A) \approx \hat{I}_B/D$ and $q_{z_A}(E) \approx D_B/D^2$, where D and D_B are the Hilbert-space dimension of the total system and subsystem B respectively, meaning

$$\hat{\rho}_d(z_A) \approx \frac{1}{D} \hat{I}_B \quad (7.45)$$

$$\text{tr}_B[\hat{\rho}_d(z_A)] \approx \frac{D_B}{D} \quad (7.46)$$

$$\text{tr}_B[\hat{\rho}_d(z_A)^2] \approx \frac{D_B}{D^2} \quad (7.47)$$

$$q_{z_A}(E) \approx \frac{1}{D} \sum_{z_B} \frac{1}{D} \approx \frac{D_B}{D^2} \quad (7.48)$$

$$\frac{1}{D_\beta} \approx \frac{D(D_B/D^2)^2}{(D_B/D)^2} \approx \frac{1}{D}, \quad (7.49)$$

leading to $D_\beta \rightarrow D$ and $D_{B,\text{eff}} \rightarrow D_B$.

Applying these infinite temperature approximations, we immediately arrive at the result of Eq. (7.2), that

$$\sigma_{\text{rel}}^2(z_A) = \frac{1}{D_B} - \frac{1}{D}. \quad (7.50)$$

Explicitly, in Fig. 7.4 in the main plot for the Rydberg, mixed field Ising model, and random unitary circuits the x-axis is D_B , while for the Bose-Hubbard system the x-axis is $D_{B,\text{eff}}$ to account for particle number conservation. Further, when calculating the scaling coefficients for all models in the inset of Fig. 7.4, we use $D_{B,\text{eff}}$ instead of D_B - this is to be maximally fair to all models in case there are slight non-infinite temperature conditions.

The prediction of Eq. (7.42) is also matched by both numerical and experimental data at finite effective temperature. In Fig. 7.8, we plot the relative variance, $\sigma_{\text{rel}}^2(z_A)$ versus either the bath dimension, D_B , or the effective bath dimension, $D_{B,\text{eff}}$ (Eq. (7.43)), for finite temperature evolution with both Rydberg (experiment and simulation) and mixed field Ising model (simulation) systems. We observe the relative variance scales inversely with the effective bath dimension, as expected.

To be explicit, for each bitstring z_A we calculate the corresponding $D_{B,\text{eff}}$ and D_β numerically, which will be different generically for every z_A . For the experimental data, we select only blockade-satisfying bitstrings as we do not have sufficiently many experimental measurements to resolve probabilities with high confidence outside this space. However, we note in general that rescaling the bath dimension naturally takes into account the Rydberg blockade, as it is essentially a finite temperature effect. In other words, while we can operate in an infinite effective temperature regime within the blockade-satisfying subspace, within the global Hilbert space the system is at effectively low temperature; this temperature is naturally accommodated via the bath dimension rescaling for each z_A .

7.4.6 Higher moments of subsystem bitstring probabilities

Eq. (7.42) shows that the relative temporal variance, i.e. the second moment of the PoP, approximately satisfies $\sigma_{\text{rel}}^2 = 1/D_B - 1/D$. This agrees with the second moment of the Erlang distribution, up to the $1/D$ term. Here, we explicitly compute the third moment, demonstrate that it is similarly consistent with the Erlang distribution and discuss leading order corrections.

Higher moments of the temporal ensemble may be evaluated [30, 71]. We find:

$$\begin{aligned} \mathbb{E}_t[p(z_A, t)^3] &= \text{tr}[\hat{\rho}_d(z_A)]^3 + 3\text{tr}[\hat{\rho}_d(z_A)]\text{tr}[\hat{\rho}_d(z_A)^2] + 2\text{tr}[\hat{\rho}_d(z_A)^3] \\ &\quad - 3\text{tr}[\hat{\rho}_d(z_A)]\text{tr}[\hat{\rho}_d^{(2)}(z_A)] - 6\text{tr}[\hat{\rho}_d^{(2)}(z_A)(\hat{I} \otimes \hat{\rho}_d(z_A))] + 4\text{tr}[\hat{\rho}_d^{(3)}(z_A)] \end{aligned} \quad (7.51)$$

where $\hat{\rho}_d^{(k)}(z_A) \equiv (\langle z_A | \otimes \hat{I}_B)^{\otimes k} \sum_E |c_E|^{2k} |E\rangle \langle E|^{\otimes k} (|z_A\rangle \otimes \hat{I}_B)^{\otimes k}$ are higher-order correction terms. We diagrammatically illustrate this expression in Fig. 7.7B. We emphasize that this expression is exact, and further approximations are often required to make headway.

The relevant quantity we wish to compare is the third cumulant κ_3

$$\kappa_3 \equiv \frac{\mathbb{E}_t[p(z_A, t)^3]}{\mathbb{E}_t[p(z_A, t)]^3} - 3 \frac{\mathbb{E}_t[p(z_A, t)^2]}{\mathbb{E}_t[p(z_A, t)]^2} + 2 \quad (7.52)$$

$$= 2 \frac{\text{tr}[\hat{\rho}_d(z_A)^3]}{\text{tr}[\hat{\rho}_d(z_A)]^3} - 6 \frac{\text{tr}[\hat{\rho}_d^{(2)}(z_A)(\hat{I} \otimes \hat{\rho}_d(z_A))]}{\text{tr}[\hat{\rho}_d(z_A)]^3} + 4 \frac{\text{tr}[\hat{\rho}_d^{(3)}(z_A)]}{\text{tr}[\hat{\rho}_d(z_A)]^3} \quad (7.53)$$

$$\approx \frac{2}{D_B^2} - \frac{6}{D_B D} + \frac{4}{D^2} \quad (7.54)$$

where in the last step we have used the same approximations above at infinite temperature, that $\hat{\rho}_d(z_A) \approx \hat{I}_B/D$ and $q_{z_A}(E) \approx D_B/D^2$ (and note that $\text{tr}(\hat{\rho}_d^{(3)}(z_A)) = \sum_E q_{z_A}(E)^3$). As with the variance, this agrees with the third cumulant of the Erlang distribution (which is simply $2/D_B^2$) up to corrections $O(1/D)$. It is also close, but not exactly equal to $2(1/D_B - 1/D)^2$, i.e. replacing $1/D_B$ with $1/D_B - 1/D$ as the parameter in the Erlang distribution.

We note that both Eqs. (7.50) and (7.54) agree with the (average) second and third cumulants of Haar random states [285, 340], i.e. for marginal distributions $p_\psi(z_A) \equiv \langle z_A | \text{tr}_B(|\psi\rangle\langle\psi|) | z_A \rangle$ of Haar random states $|\psi\rangle$,

$$\mathbb{E}_{|\psi\rangle \sim \text{Haar}} [D_A^2 p_\psi^2(z_A)] - 1 = \frac{D + D_A}{D + 1} - 1 \approx \frac{1}{D_B} - \frac{1}{D} \quad (7.55)$$

$$\mathbb{E}_{|\psi\rangle \sim \text{Haar}} [D_A^3 p_\psi^3(z_A) - 3D_A^2 p_\psi^2(z_A) + 2] = \frac{D^3 + 3D^2 D_A + 2DD_A^2}{D(D+1)(D+2)} - 3 \frac{D + D_A}{D + 1} + 2 \quad (7.56)$$

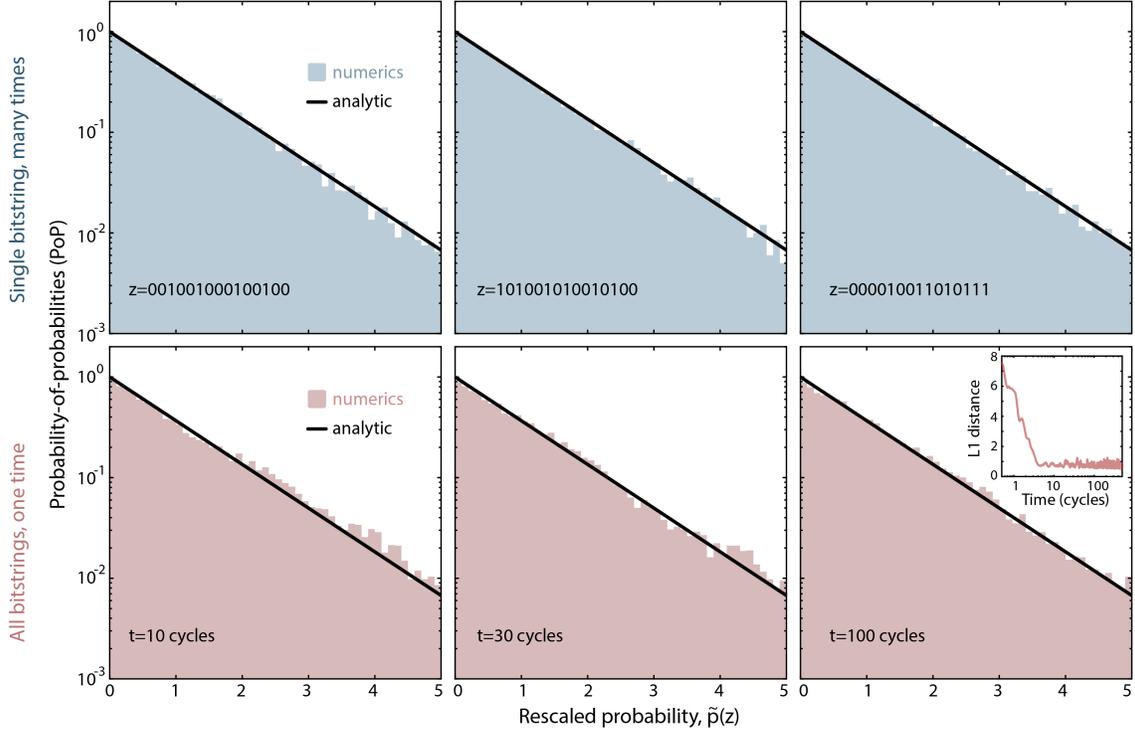


Figure 7.9: **Observation of Porter-Thomas statistics on the level of a single bitstring or time.** Probability-of-probabilities (PoP) distributions obtained from numerical simulations for an $N = 15$ Rydberg atom array. Top: PoP distributions for the probability to observe a single bitstring, aggregated over evolution times ranging from 0 to 500 cycles, and normalized by the average temporal probability. Bottom: PoP distributions aggregated over all 2^{15} bitstrings at a single fixed time, with probabilities normalized by their average temporal probability. In all cases, the observed PoP distributions agree with the exponential Porter-Thomas distribution, as is expected from theoretical arguments related to Hilbert-space ergodicity [71]. In the case of averaging over all bitstrings at a fixed time, this agreement is observed to set in past ≈ 4 cycles, as quantified by the L1-norm between the numerical and analytic PoP as a function of time (inset). Note that in all plots we consider not just the lowest energy non-blockaded sector, but the entire Hilbert space.

$$\approx \frac{2}{D_B^2} - \frac{6}{D_B D} + \frac{4}{D^2} \quad (7.57)$$

We next note that the substitutions $\hat{\rho}_d(z_A) \approx \hat{I}_B/D$ and $q_{z_A}(E) \approx D_B/D^2$ leading up to Eq. (7.54) are heuristic in nature and they may not appropriately capture the relevant higher-order fluctuations to the desired level of precision (here, up to $1/D^2$). Nevertheless, they appear to give the correct answer; we have verified that Eqs. (7.50) and (7.54) agree with a more careful calculation which treats the eigenbasis $\{|E\rangle\}$ as a random basis, and averages over such bases with an integral over Haar-random unitaries.

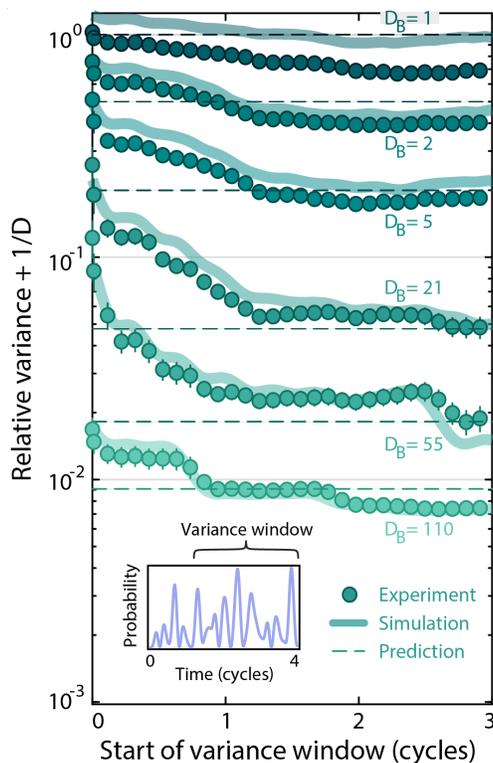


Figure 7.10: **Time-resolved saturation of the temporal variance.** We plot $\sigma_{\text{rel}}^2 + 1/D$ for both experiment and numerical simulations evolved at infinite temperature, for various choices of bath dimension, D_B . The variance is calculated in a temporal window of varying length, where the endpoint is fixed at 4 Rabi cycles of evolution, the latest time measured experimentally, while the start of the window is varied (inset). We see for both the numerics and the experiment the variance settles to the prediction of $1/D_B$ for start times of ≈ 1.5 cycles. We note that when simulating out to many hundreds of cycles – to emulate the infinite time limit – the convergence is observed without excluding any early time transient behavior (not shown), but we do not have the experimental resolution to evolve to such late times without incurring fidelity loss.

Finally, we note that in all approaches discussed above, the leading $O(1/D_B^{k-1})$ terms of all k -th cumulants agree with those of the Erlang distribution.

7.4.7 Bitstrings vs time, and time-dependence of the PoP

So far, we have shown the global PoP forms the Porter-Thomas distribution when aggregated over many measurement times. However, under Hilbert-space ergodicity, Refs. [30, 71] state the Porter-Thomas distribution should also be apparent when forming the PoP over many choices of bitstrings at a fixed measurement time.

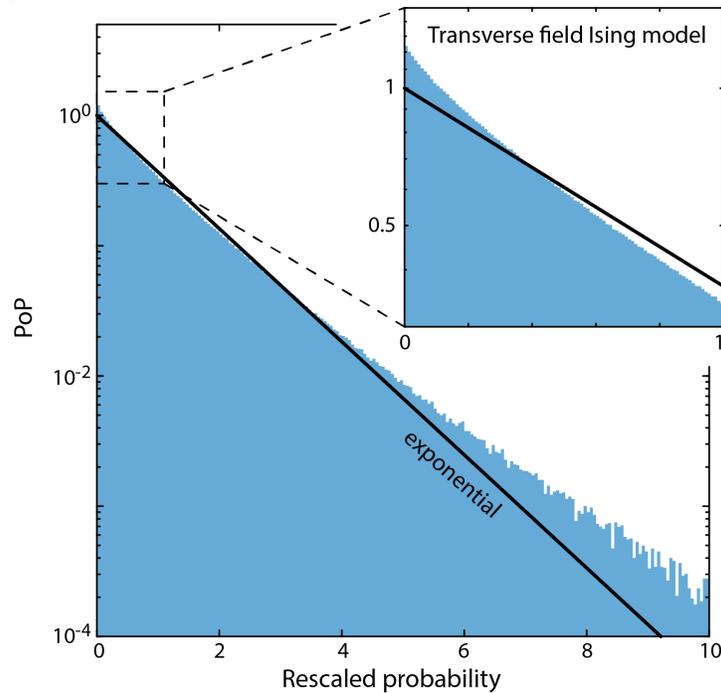


Figure 7.11: **Non-Porter-Thomas behavior of the transverse field Ising model.** We evolve an $N = 12$ atom system under the transverse field Ising model out to several hundred cycles, as in the inset of Fig. 7.4, and aggregate all global bitstring probabilities over all bitstrings and at different times (to obtain better statistics), normalized by their respective long time averages. The resulting PoP shows non-exponential and non-universal signatures. Accounting for such effects may allow us to better describe models not exhibiting Hilbert-space ergodicity.

In Fig. 7.9 we present numerical evidence for this by simulating the Rydberg Hamiltonian with no decoherence out to several hundred cycles of evolution for $N = 15$ atoms. We then form the global PoP either by a) aggregating normalized probabilities across many times for a single bitstring, or b) aggregating normalized probabilities over many bitstrings at a fixed time. Explicitly, in both cases the normalization is done by dividing probabilities by their respective long-time averages. Note that typically for the Rydberg systems we only consider the non-blockaded Hilbert space as it is difficult to gather sufficient experimental statistics to reconstruct bitstring probabilities in the higher-energy sectors, but in Fig. 7.9 we consider all bitstrings, not just those drawn from the non-blockaded subspace.

An important aspect visible in the inset of Fig. 7.9 is that the fixed-time PoP converges close to the Porter-Thomas distribution over a few cycles of evolution. In Fig. 7.2 – where we observe the experimental global PoP is roughly consistent with the Porter-Thomas

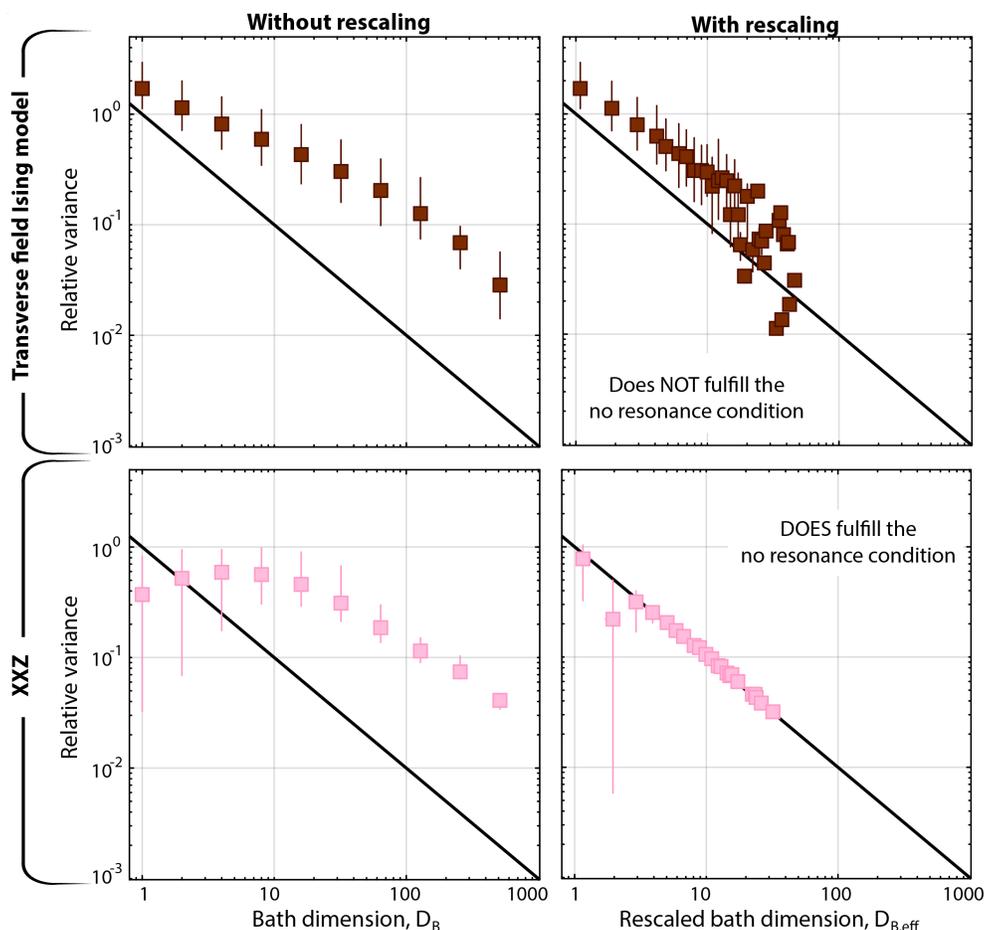


Figure 7.12: **Integrable models exhibiting Hilbert-space ergodicity.** While some integrable models (e.g. the transverse field Ising model) do not fulfill the no-resonance conditions, and thus do not exhibit HSE, some integrable models (e.g. the XXZ model) do. Thus, we expect, and observe that temporal fluctuations of the XXZ model obey the universal scaling of Eq. (7.2), similar to models studied in Fig. 7.4.

distribution – we aggregated over both bitstring index and measurement times from 1.5 to 4 cycles of evolution to gather sufficient statistics. Evolving out to longer times experimentally modifies the PoP distribution due to interaction with the external environment, as discussed at length. To verify we study sufficiently late dynamics in the experiment, we consider σ_{rel}^2 calculated in a temporal window of varying length, but fixed endpoint at 4 Rabi cycles of evolution (Fig. 7.4), the latest time measured experimentally. We find σ_{rel}^2 appears roughly converged for a start time of ≈ 1.5 cycles, for which we observe both the experimental and numerical σ_{rel}^2 closely matches the $1/D_B - 1/D$ prediction. We note some disagreement for the experimental data at $D_B = 1$, which again we attribute to early signatures of decoherence.

7.4.8 Failure cases

As shown in the inset of Fig. 7.4, not all models follow our analytic predictions, in particular those that do not fulfill the *no-resonance conditions* such as the transverse field Ising model and systems of non-interacting spins. As a further example of this failure we plot the PoP aggregated over global bitstring probabilities from many times for the transverse field Ising model (Fig. 7.11). Probabilities are normalized by their long time average, and are measured out to several hundred cycles. We observe clear non-exponential signatures in the PoP, which then lead to the non-universal features seen in Fig. 7.4. However, this is not to imply all integrable models behave in this way; some integrable models, like the XXZ Hamiltonian, do fulfill the no-resonance conditions, and thus do exhibit Hilbert-space ergodicity, as typified in Fig. 7.12.

7.5 GENERALIZING TO HYPOEXPONENTIAL DISTRIBUTIONS

In Section 7.4, we formally derived the emergence of the Erlang distribution from results in Hilbert-space ergodicity for closed system dynamics. Using the intuition gained from that derivation, in this section we build models of how the probability-of-probabilities (PoP) should behave under a variety of conditions, including open system dynamics, explicitly comparing them to numerical simulations. We find that the PoP qualitatively and quantitatively is changed in the presence of different types of noise, which we used both to discriminate between noise channels as well as learn their relative strengths.

We first give an overview of how the exponential, or Porter-Thomas distribution emerges from ergodic unitary dynamics. We then introduce hypoexponential distributions, which parameterize all the PoP distributions studied in this work, including the Erlang distribution which describes closed system dynamics at infinite temperature. For closed system dynamics at finite temperature, we describe the hypoexponential weights and discuss the conditions on the bath size under which these weights are the eigenvalues of the system (or bath) reduced density matrix.

We then move on to the case where our system is coupled to an external environment, describing our models for local errors, global depolarizing noise, Gaussian coherent errors, and a mixture of local and Gaussian coherent errors.

In this section, all probability distribution functions (PDFs) will be supported on the non-negative real line $x \geq 0$.

7.5.1 Porter-Thomas from random distributions

Our basic building block will be the exponential distribution $\text{Exp}(\lambda)$, with PDF

$$P_{\text{Exp}}(x, \lambda)dx = \lambda \exp(-\lambda x)dx \quad (7.58)$$

We will refer to variables which follow the exponential distribution with $\lambda = 1$ (or, *Porter-Thomas distribution*) as “standard exponential variables.”

The Porter-Thomas distribution is a universal feature of random quantum states [275, 289, 297]. To see this, note that the probability distributions $p(z) = |\langle z|\Psi\rangle|^2$ obtained from Haar-random states are random probability distributions $p(z)$. To be precise, such distributions are drawn from the *flat Dirichlet distribution*, which is a distribution over the space of probability distributions (the probability simplex). Such random distributions can be generated by i.i.d. sampling each probability value $p(z) \sim \text{Exp}(1/D)$, and then normalizing all entries such that $\sum_z p(z) = 1$. This final step introduces weak correlations $\text{cov}[p(z), p(z')] = \delta_{z,z'}/2D - 1/2D^2$, which vanish in the limit of large Hilbert-space dimension D .

We are primarily interested in the marginal distributions $P[p]$ of individual probability values, namely, the distribution of the values of $p(z)$. This is known to be the *Beta distribution* $\text{Beta}(1, D-1)$ [291, 355], which converges to $\text{Exp}(1/D)$ in the large system limit $D \rightarrow \infty$. Crucially, the marginal distribution $P[p(z)]$ can be observed from a single *typical* distribution $p(z)$, and one does not have to sample many Dirichlet random distributions. Since the entries $p(z)$ of a typical distribution $\{p(z)\}$ are very weakly correlated, the histogram of their values, or probability-of-probabilities (PoP), is also approximately exponentially distributed. In this context, the exponential distribution is referred to as the Porter-Thomas distribution.

In the above discussion, we have motivated the Porter-Thomas distribution through Haar-random states. However, the Porter-Thomas distribution is present in states which can be practically produced. Specifically, it is present in the states output from random unitary circuits (RUCs) [275]. It is convenient to rescale the probability values as $Dp(z)$, which follow the standard exponential distribution (i.e. with mean 1).

Furthermore, as shown in Section 7.4.2, the Porter-Thomas distribution appears in states obtained from unitary Hamiltonian evolution [30], i.e. $|\Psi(t)\rangle = \exp(-i\hat{H}t)|\Psi_0\rangle$ for an arbitrary initial state $|\Psi_0\rangle$. We prove [71] this in the limit of exponentially long times, and observe that in practice the required time is small, only growing weakly with system size [30]. Specifically, we consider the probabilities $p(z, t) \equiv |\langle z|\Psi(t)\rangle|^2$ and their time-average values $p_{\text{avg}}(z) \equiv \mathbb{E}_t[p(z, t)]$. Hamiltonian dynamics differs from RUC dynamics

by energy conservation, which introduces non-trivial $p_{\text{avg}}(z)$. The probabilities $p(z, t)$ fluctuate about their average value, but the rescaled probabilities $\tilde{p}(z) \equiv p(z)/p_{\text{avg}}(z)$ follow the Porter-Thomas distribution.

7.5.2 Hypoexponential distributions

As described, the hypoexponential distribution is the probability distribution of a weighted sum of independent standard exponential variables. For a given set of weights ω_i (which are positive and sum to 1), and associated i.i.d. Dirichlet probability distributions $\{p_i(z)\}$, the overall distribution is given by the weighted sum:

$$p(z) = \sum_i \omega_i p_i(z). \quad (7.59)$$

When the number of weights is small (i.e. for a small bath dimension), the weights can be related to the eigenvalues of the density matrix $\hat{\rho}$. Their properties may be inferred from each other (Section 7.7.1). For example, the state purity is related to the second moment of the weights $\text{tr}(\hat{\rho}^2) = \sum_i \omega_i^2$, which we utilize heavily in Section 7.7.

The corresponding PoP distribution is given by $\text{Hypo}(\{\omega_i\})$ (Note that the hypoexponential distribution is conventionally parameterized in terms of the inverse weights ω_i^{-1} , which are known as *rates*). When the ω_i are distinct, the hypoexponential distribution has probability distribution function (PDF) [347]

$$P_{\text{Hypo}}(x) = \bigotimes_i P_{\text{Exp}}(x/\omega_i) = \sum_i \omega_i^{-1} \exp(-x/\omega_i) \prod_{j \neq i} \frac{\omega_j}{\omega_j - \omega_i}, \quad (7.60)$$

where \otimes denotes convolution, an operation of functions defined by $(f \otimes g)(x) = \int dy f(y)g(x-y)$. Eq. (7.60) does not hold when there are degenerate ω_i , which is often the case for our purposes. Hence, we will rarely directly use Eq. (7.60) (see Ref. [347] for an expression valid for degenerate ω_i).

Different types of noise give different weights ω_i . Below, we state types of noise we consider, their weights ω_i and their resultant PoP distributions.

7.5.3 Intrinsic baths

We first consider the PoP distribution in the context of closed system quantum thermalization: a subsystem effectively thermalizes due to its entanglement with the rest of the system, which serves as a bath.

Infinite temperature

First, we let the bath be at infinite effective temperature. The quantities of interest are the subsystem bitstring probabilities $p(z_A)$, which are marginal probabilities $p(z_A) = \sum_{z_B} p(z_A \cdot z_B)$. As discussed, we find that the PoP of the marginal probabilities are described by Erlang distributions. In Section 7.4, we gave a rigorous derivation of this, including small deviations from the perfect Erlang distribution. Here, we give an intuitive, heuristic derivation.

Consider an infinite temperature state in the absence of symmetries for simplicity. We take the average probabilities $p_{\text{avg}}(z)$ to be trivially $1/D$. Therefore, the joint probabilities $p(z_A \cdot z_B)$ are essentially i.i.d. exponential variables, and the (rescaled) sum $D_A p(z_A)$ is the average of D_B i.i.d. standard exponential variables, i.e. $\omega_i = 1/D_B$ for $i = \{1, 2, \dots, D_B\}$. This special case of the hypoexponential distribution is precisely the *Erlang distribution*, parameterized by a positive integer k (here $k = D_B$). Because the weights are degenerate, its PDF cannot be directly obtained from Eq. (7.60), but can be independently derived as

$$P_{\text{Erlang}}(x; k) = \left(\bigotimes_{i=1}^k \right) P_{\text{Exp}}(kx) = \frac{k^k x^{k-1}}{(k-1)!} \exp(-kx), \quad (7.61)$$

where in our case of interest the variable $x = D_A p(z_A)$, and $k = D_B$.

Intrinsic bath at finite temperature

Having established the PoP distribution for an intrinsic bath at infinite effective temperature, a natural next step is to investigate the PoP distribution at finite effective temperatures. Specifically, we consider the following setting: a system is partitioned into a larger subsystem A and smaller one B , which we treat as the “bath.”

We first restrict ourselves to a special type of measurement basis z_A , whose conditions are explained later. Here, we claim that when the bath B is small, the PoP of the bitstrings on A is the hypoexponential distribution with coefficients λ_i^{-1} , where $\{\lambda_i\}_{i=1}^{D_B}$ are the eigenvalues of the reduced density matrix $\hat{\rho}_A$ (or $\hat{\rho}_B$). That is,

$$\tilde{p}(z_A) \equiv p(z_A)/p_{\text{avg}}(z_A) \sim \text{Hypo}(\{\lambda_i\}), \quad (7.62)$$

where $p(z_A) = \sum_{z_B} p(z_A \cdot z_B)$ and $p_{\text{avg}}(z_A) = \sum_{z_B} p_{\text{avg}}(z_A \cdot z_B)$ are the marginal distributions of z_A .

This result may be understood through the lens of *projected ensembles* [29, 70]. These are the ensembles of states defined on a subsystem Hilbert space (here, \mathcal{H}_B) by measuring a global state $|\Psi\rangle$ in its complement (A). For concreteness, we let $|\Psi\rangle = |\Psi(t)\rangle$ be generated by many-body dynamics for time t . Each of the exponentially many measurement outcomes

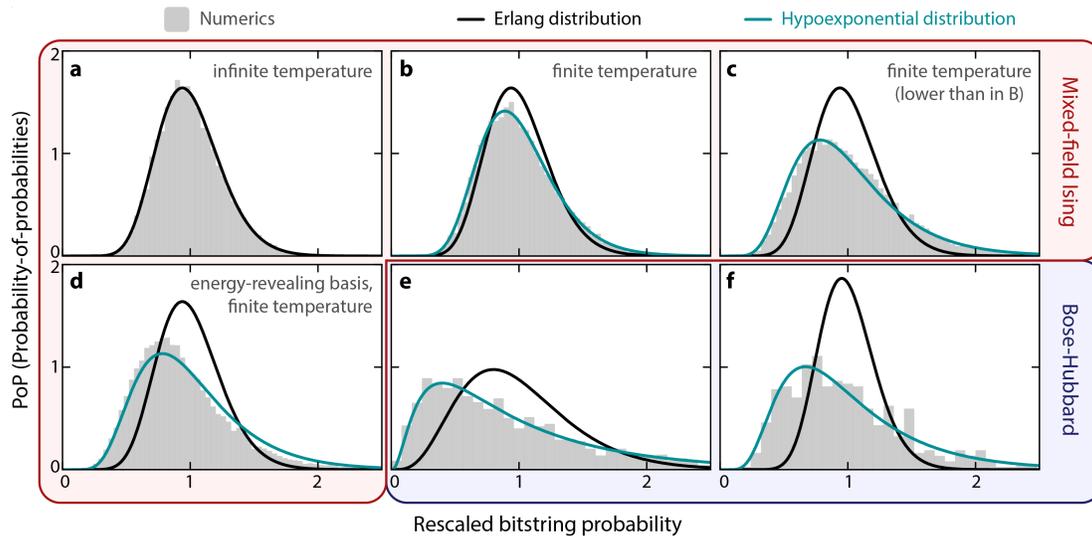


Figure 7.13: **Hypoexponential distributions for finite temperature closed system dynamics.** In (a-d), we plot the PoP of the conditional probabilities of $\tilde{p}(z_A) = p(z_A)/p_{\text{avg}}(z_A)$ formed by states time-evolved by the mixed field Ising model (MFIM). Here, the system is a one-dimensional $N = 18$ open chain and the subsystem B is four sites at the center of the chain. (abc) The temperature is tuned by rotating the initial state $|\Psi_0\rangle = \left[\exp(i\theta \hat{S}_x) |0\rangle \right]^{\otimes N}$ at various angles $\theta = 0, -0.6, -1.2$. In abc, B is measured in the $\{z_B\}$ basis. Only at infinite temperature does the PoP agree with the Erlang distribution (green), while at finite temperatures, the PoP agrees with the hypoexponential distribution [Eq. (7.62)] instead. In d, B is measured in the energy-revealing $\{x_B\}$ basis, showing small deviations from the hypoexponential distribution. In ef, we form the PoP from the state $|1, 1 \dots 1\rangle$ evolved by the Bose-Hubbard model [Eq. (7.7)] with $M = 10$ sites and $N = 10$ particles, measured in the particle number basis, and with a subsystem of 3 sites. Due to the number-conservation of the evolution, we restrict to the measurement outcomes where z_B has e) 4 and f) 5 particles, and form the symmetry-resolved density matrix accordingly. The PoPs agree better with the hypoexponential prediction instead of the Erlang distribution with the appropriate symmetry-resolved subsystem dimension.

z_A projects $|\Psi\rangle$ into a distinct pure state $|\Psi(z_A)\rangle \in \mathcal{H}_B$. The statistical properties of projected ensembles translate to properties of the PoP.

In companion work [71], we study projected ensembles at finite temperature, relevant to our setting here. When two conditions are satisfied, $\hat{\rho}_d(z_A)$ is approximately the reduced density matrix $\hat{\rho}_B(t) \equiv \text{tr}_A [|\Psi(t)\rangle\langle\Psi(t)|]$, which gives the stated claim Eq. (7.62). These two conditions are (i) that the measurement basis should be “energy non-revealing,” and (ii) the bath size must be sufficiently small.

Effect of the measurement basis — The result Eq. (7.62) applies to specific measurement bases $|z_A\rangle$ that are “energy non-revealing,” i.e. do not have significant correlations with the Hamiltonian that generated the state $|\Psi(t)\rangle$ (and which defines the notion of temperature). One example of an “energy non-revealing basis” is the Z_A basis for the mixed field Ising model with fields and couplings in the X and Y directions, where the measurement outcomes z_A are uncorrelated with the energy density. For other bases, such as the X_A basis for the mixed field Ising model discussed above, such correlations may exist. Specifically, due to energy conservation of the overall state, measurement outcomes with higher values of energy (on subsystem A) on average lead to projected states with lower energy (on subsystem B), and vice-versa. Therefore the projected ensemble $\{|\Psi(z_A)\rangle\}$ may not be taken to be samples from *the same distribution*, independent of z_A . Among other consequences [71], the simple prediction Eq. (7.62) does not hold in this setting.

Effect of bath size — When the measurement basis is energy non-revealing, we may conclude that the *time-averaged* reduced density matrix $\mathbb{E}_t[\hat{\rho}_B(t)]$ is proportional to any of the $\hat{\rho}_d(z_A)$. However, $\hat{\rho}_B(t)$ has fluctuations over time, and therefore the eigenvalues of $\mathbb{E}_t[\hat{\rho}_B(t)]$ (the weights of the hypoexponential PoP) may not agree with the eigenvalues of $\hat{\rho}_B(t)$. To see this, it is instructive to consider the case of infinite temperature, where we model the global state as a Haar-random state. Indeed, for Haar-random states, the eigenvalues of the reduced state follow the so-called Marchenko-Pastur distribution [356, 357], which is supported on the interval $[D_B^{-1}(1 - \sqrt{D_B/D_A})^2, D_B^{-1}(1 + \sqrt{D_B/D_A})^2]$ and is controlled by the parameter $D_B/D_A = 2^{2|B|-N}$. Meanwhile, the time-averaged reduced state is equal to \hat{I}_B/D_B , with all eigenvalues equal to $1/D_B$. When $|B| = N/2$, the second moment of the Marchenko-Pastur distribution is $\sum_j \lambda_j^2 \approx 2/D_B$, while the flat spectral distribution has second moment $1/D_B$, a clear disagreement. This feature is not particular to our results, and instead reflects the conditions under which the time-dependent reduced density matrix may be approximated by a time-independent thermal state. Thus, the size of the bath should be sufficiently small; specifically, $|B|$ should be a constant amount less from half the system size $N/2$.

In Fig. 7.13, we confirm some of our predictions. In Fig. 7.13ABC we plot the PoP of a state time-evolved under the MFIM [Eq. (7.5)], with a total system size of $N = 18$ qubits, and a bath size of $N_B = 4$ qubits. We tune the effective temperature by rotating the initial state $|\Psi_0\rangle = \left[\exp(i\theta\hat{S}_x)|0\rangle\right]^{\otimes N}$. At infinite temperature (A, $\theta = 0$), the PoP is well described by the Erlang distribution. Tuning through finite temperature (BC, $\theta = -0.6, -1.2$), the PoP is instead well described by the hypoexponential distribution [Eq. (7.62)]. Measuring in the energy-revealing $\{x_B\}$ basis (D), the PoP shows small deviations from the Eq. (7.62). To illustrate the ubiquity of our prediction, we illustrate this with a finite temperature state evolved under the Bose-Hubbard model. In the presence of number conservation $\langle\hat{N}\rangle = N_{\text{tot}}$, we must form the projected ensemble by restricting to measurement outcomes with fixed number \hat{N}_A , and forming the according symmetry resolved density matrix $\hat{\rho}_B(\hat{N}_A) \propto \sum_{\{z_A|n(z_A)=\hat{N}_A\}} (\langle z_A| \otimes \hat{I}_B) |\Psi(t)\rangle \langle\Psi(t)| (|z_A\rangle \otimes \hat{I}_B)$, as detailed in Refs. [70, 71].

7.5.4 Extrinsic baths

We now consider the appearance of the hypoexponential under coupling to an extrinsic bath, i.e. with weights controlled by a given noise model.

Local errors

First, we consider the case of local errors. As illustrated in Fig. 7.5, our model for the PoP obtained from local errors with k locations is that it is the convolution of one exponential variable (speckle pattern) with large weight F (the ‘‘ideal trajectory’’), and k exponential variables with small weights $(1 - F)/k$. In this case, the PDF is given by

$$P_{\text{Local Err.}}(x; F, k) = P_{\text{Exp}}(x/F) \otimes P_{\text{Erlang}}(x/(1 - F); k) \quad (7.63)$$

$$= \left(\frac{k}{1 - F^{-1} + k}\right)^k \gamma\left(k, \frac{F - 1 + kF}{F(1 - F)}x\right) \frac{e^{-x/F}}{F(k - 1)!}, \quad (7.64)$$

where $\gamma(k, x) \equiv \int_0^x z^{k-1} e^{-z} dz$ is the lower incomplete Gamma function. In Fig. 7.14ab we plot these PDFs for various values of k and F .

This model holds for any type of local error. In Fig. 7.5b, we illustrate the PoP with amplitude damping errors, i.e. spontaneous decay of the $|1\rangle$ state to $|0\rangle$, while in Fig. 7.14 we show it for local Pauli noise. In practice, however, the various assumptions made for the local noise channel may not hold, for instance because errors are not necessarily completely orthogonal and because multiple errors can occur in a single shot of the experiment. These largely result in a value of k which is not simply the spacetime volume of the circuit. To

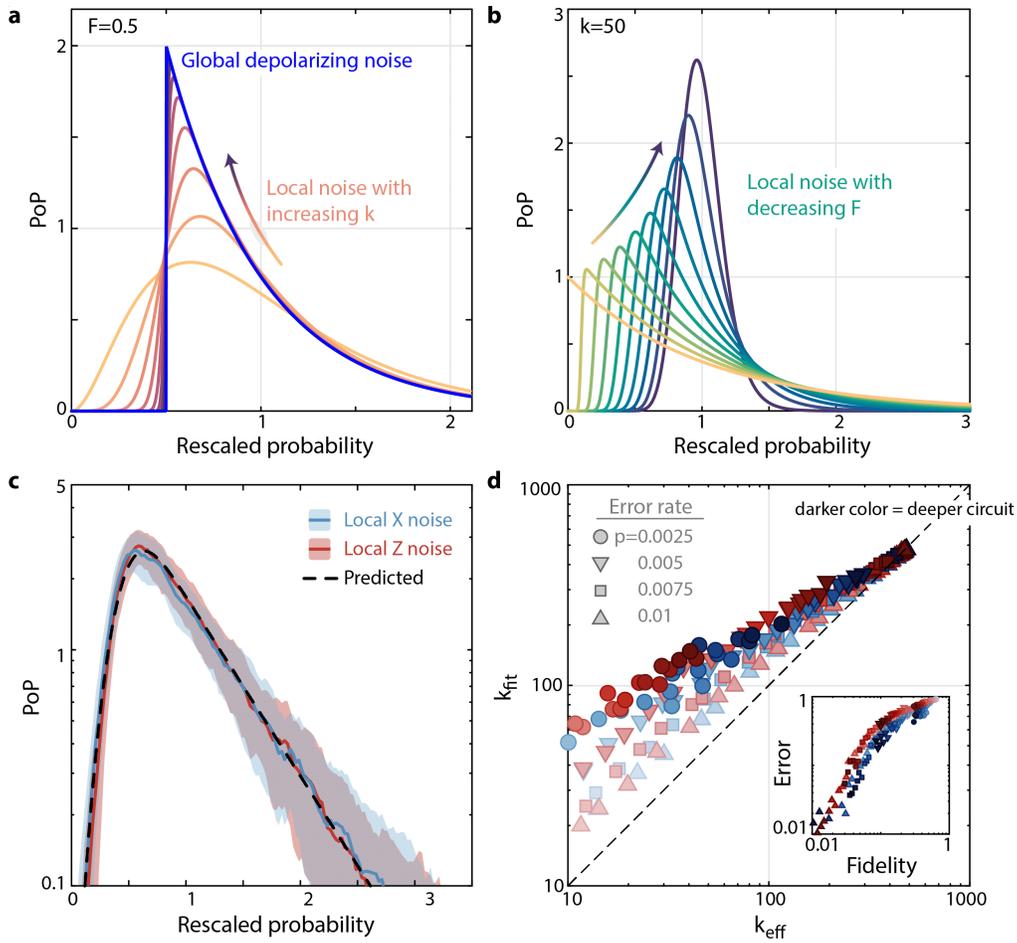


Figure 7.14: Probability-of-probabilities (PoP) for local noise. **a.** PoP distributions for local noise (see Eq. (7.63)) with $F = 0.5$ and k varying from 2 to 10000 logarithmically. As k increases, the PoP distribution converges to the prediction for global depolarizing noise (see Eq. (7.66)) at the same fidelity, consistent with Ref. [348]. **b.** PoP distribution for local noise with $k = 50$ and F varying from 1 to 0.1. **c.** PoP distribution for either local Pauli-X errors or Pauli-Z errors, both of which are roughly equally well described by the same effective k value. **d.** Fitted value of k versus value predicted from Eq. (7.65) for noisy simulation with either Pauli-X or Pauli-Z errors evaluated for varying circuit depths. Inset: The fitted value agrees best with the predicted value in the low fidelity regime, as quantified by the relative error between k_{eff} and k_{fit} . In **c-e**, results are aggregated over 10 random circuit instances, blue points represent a local Pauli-X error model, and red points represent a local Pauli-Z error model.

address this, in practice we use an effective predicted k value defined as

$$k_{\text{eff}} = \frac{(1 - F)^2}{\kappa_2 - F^2}, \quad (7.65)$$

where κ_2 is the second cumulant of the experimentally measured PoP distribution (as

discussed in greater detail in Section 7.7, see Table 7.1). Using Eq. (7.65) we can predict the k value for a given measured PoP distribution at a given fidelity. In Fig. 7.14c we show the good agreement between Eq. (7.63) and numerical simulations using an effective k found from Eq. (7.65). Further, in Fig. 7.14d we compare k_{eff} from Eq. (7.65) against k value of best fit for various error rates and evolution depths, and generally find the two are in good agreement, particularly in the low fidelity regime (Fig. 7.14d, inset).

We stress that this model is limited. As an example, it stipulates that the weights of the non-ideal trajectories be equal. This is not true if more than one error can occur during the dynamics: the probability of the trajectory decreases exponentially with the number of errors, while the number of such distinct trajectories increases exponentially. Furthermore, even the weights of single-error events may not be uniform. For example, the probability of amplitude damping errors depend on the magnetization of the state at each layer, which may change during the dynamics. These effects may lead to small deviations between our model and PoPs generated by realistic many-body dynamics, which may affect high-precision tasks such as determining the fraction of coherent-to-local error in Fig. 7.6bc. However, such limitations may be solved with straightforward model-level refinements (such as explicitly adjusting weights to account for multiple error occurrences).

Extrinsic bath: Global depolarizing noise

Global depolarizing noise is a commonly used toy model for the effects of noise on quantum states. This channel maps states $\hat{\rho} \mapsto F\hat{\rho} + (1 - F)\hat{I}/D$, where \hat{I} is the identity and F is a parameter which is approximately the fidelity of the channel. Under this channel, the bitstring distribution $p(z) \mapsto Fp(z) + (1 - F)/D$ and hence the PoP distribution of $x = Dp(z)$ is shifted to the right and rescaled by $1/F$, with PDF

$$P_{\text{GD}}(x; F) = \begin{cases} 0 & \text{for } 0 \leq x < 1 - F \\ F^{-1} \exp\left(-\frac{x-(1-F)}{F}\right) & \text{for } x \geq 1 - F \end{cases} \quad (7.66)$$

As seen in Fig. 7.14, $P_{\text{Local Err.}}(x; F, k) \rightarrow P_{\text{GD}}(x; F)$ as $k \rightarrow \infty$, consistent with Ref. [348], which states that the output distribution of random circuit sampling under local noise converges to that of “global white noise” when scaling the circuit spacetime volume.

Extrinsic bath: Gaussian coherent errors

Next, we discuss the case of Gaussian coherent errors, which describes shot-to-shot variation of a global parameter.

For a circuit model, we model this error by systematically under- or over-rotating every two-qubit gate by a multiplicative, Gaussian-distributed factor. For Hamiltonian

evolution, a global Hamiltonian parameter, such as the Rabi frequency Ω , is taken to be a Gaussian random variable. Such a noise channel is linked to shot-to-shot parameter fluctuations observed in Hamiltonian evolution [31], and certain models of control errors. As found in Ref. [31], evolution for a time t under these parameter values produces a continuous family of pure states $|\Psi(t, \Omega)\rangle$ which mutually overlap as $|\langle\Psi(t, \Omega)|\Psi(t, \Omega')\rangle|^2 \propto \exp(-N(\Omega - \Omega')^2 t^2 / 2\sigma_\Omega^2)$. When Ω is normally distributed, i.e. with $P(\Omega) = \exp(-(\Omega - \Omega_0)^2 / (2\sigma'^2)) / \sqrt{2\pi\sigma'^2}$ for some σ' , we obtain a normally-distributed mixture $\hat{\rho}(t) = \int d\Omega P(\Omega) |\Psi(t, \Omega)\rangle \langle\Psi(t, \Omega)|$.

Our model for the PoP converts this continuous mixture of states into a discrete sum of probability distributions, with discrete weights ω_i . It is natural to let the weights follow a Gaussian profile when σ_F is sufficiently large (see below for deviations)

$$\omega_i^{\text{coh.}} = \frac{1}{\mathcal{Z}(\sigma_F)} \exp\left[-\frac{i^2}{2\sigma_F^2}\right], \quad (7.67)$$

where i ranges from $-\infty$ to ∞ . These weights are parameterized by the width σ_F , and $\mathcal{Z}(\sigma_F) \equiv \sum_{i=-\infty}^{\infty} \exp[-i^2/(2\sigma_F^2)]$ is a normalization constant. This sum is equal to the Jacobi theta function [358] $\theta_3(\exp[-1/(2\sigma_F^2)])$, which is well approximated by $\sqrt{2\pi}\sigma_F$ when $\sigma_F \gg 1$, but approaches 1 as $\sigma_F \rightarrow 0$. Incidentally, since $\omega_0 = 1/\mathcal{Z}(\sigma_F)$ has largest height, it corresponds to the fidelity F , i.e. the maximum overlap $\langle\Psi|\hat{\rho}|\Psi\rangle$ for any pure state $|\Psi\rangle$. Therefore, we may parameterize the PoP in terms of F instead, by solving for σ_F such that $\theta_3(\exp[-1/(2\sigma_F^2)]) = F^{-1}$.

A large fidelity corresponds to a small σ , and our model of discrete Gaussian weights shows quantitative differences in this regime. As an illustration, the second moment of the PoP can be computed analytically, and differs from the Gaussian hypoexponential model at small σ (Section 7.7 and Fig. 7.18C). This may be addressed by a non-Gaussian profile of weights in the narrow limit.

As a technical point, the weights $\omega_{-i}^{\text{coh.}} = \omega_i^{\text{coh.}}$ are degenerate and therefore we must modify Eq. (7.60) to

$$P_{\text{Gauss. Coh.}}(x; \{\omega_i\}) = P_{\text{Exp}}(x/\omega_0) \circledast \left(\bigotimes_{i=1}^{\infty} P_{\text{Erlang}}(x/\omega_i; 2) \right) \quad (7.68)$$

$$= \omega_0^{-1} e^{-x/\omega_0} \prod_{j>0} \frac{\omega_0^2}{(\omega_0 - \omega_j)^2} \quad (7.69)$$

$$+ \sum_{i=1}^{\infty} \left[\omega_i^{-2} \left(x - \sum_{\substack{j \in \mathbb{Z} \\ j \neq i}} \frac{\omega_i \omega_j}{\omega_i - \omega_j} \right) e^{-x/\omega_i} \frac{\omega_i}{\omega_i - \omega_0} \prod_{\substack{j \in \mathbb{Z} \\ j \neq i}} \frac{\omega_0}{\omega_0 - \omega_j} \right]$$

Extrinsic bath: Mixture of local and global coherent errors

Finally, we turn to modeling the effects of both local and global coherent errors simultaneously. We take the errors to be independent of each other, and the weights are simply products of the local and global weights $\omega_{ij} = \omega_i^{\text{loc.}} \times \omega_j^{\text{coh.}}$. We further notice that in numerical simulations of such a mixture of errors, the fidelity factorizes into the fidelities arising from incoherent and coherent errors (Fig. 7.18A), i.e. $F = F_{\text{coh.}} \times F_{\text{loc.}}$, i.e. $\log(F_{\text{coh.}}) + \log(F_{\text{loc.}}) = \log(F)$. This allows us to define the “fraction” of coherent errors as $f = \log(F_{\text{coh.}})/\log(F)$. We note that such a factorization is also approximately apparent for our *ab initio* Rydberg error model [31].

For completeness, we outline our algorithm used in Fig. 7.6bc to generate the PoP arising from a mixture of local and global errors based on our model, with a finite number of samples. The effects of a finite number of samples are detailed on a later section. We use this algorithm to generate the local and global coherent finite-sample PoPs in Fig. 7.6a, as well as to provide a reference distribution to calculate the χ^2 -distances in Fig. 7.6bc. Note that in the case of RUCs, the average probability is $p_{\text{avg}}(z) = 1/D$, and the rescaled empirical probabilities will only take finite values jD/M for positive integer j .

7.5.5 Other noise sources

We briefly remark that our analysis here is by no means complete. While we have analyzed commonly considered noise sources, other sources such as qubit leakage and readout error will yield different PoP distributions. For example, readout error may be treated in a similar fashion to local error. However, here there are only N possible error locations, and as such the parameter k may not be taken to be large. This is particularly notable for our moment analysis in Eq. (7.91). We leave a detailed study of this, and other noise models, to future work. However, we expect that in general accounting for novel noise models will simply amount to judicious corresponding choice of the hypoexponential weight vector.

7.6 EFFECTS OF FINITE SAMPLING

It is important to understand the effect of a finite number of experimental samples on the probability-of-probabilities discussed in this work. Finite sampling must be carefully treated since its effects on the PoP can be drastic: the simplest estimation of the PoP does not only introduce statistical noise, but also systematic bias. Here we first discuss how experimental and model PoPs are compared, then describe how the model PoP is affected by finite sampling by effectively smearing out the distribution. Finally, we suggest a method

Algorithm 1: Numerical algorithm to generate model PoP under both local and global coherent noise

Data: Fidelity F , fraction f , number of samples M , Hilbert-space dimension D , number of trajectories k , averaged probabilities $p_{\text{avg}}(z)$, and bin width $\Delta\tilde{p}$

Result: Empirical PoP \vec{n}

$F_{\text{loc.}} \leftarrow F^f$;

$F_{\text{coh.}} \leftarrow F^{1-f}$;

$\sigma_F \leftarrow \text{Solution of } \theta_3(\exp[-1/(2\sigma_F^2)]) = F_{\text{coh.}}^{-1}$;

$n(\sigma_F) \leftarrow F_{\text{coh.}}$;

$p(z) \leftarrow 0$;

for j in $\{-\lceil 5\sigma_F \rceil, -\lceil 5\sigma_F \rceil + 1, \dots, \lceil 5\sigma_F \rceil\}$; /* Infinite sum truncated at $\pm \lceil 5\sigma_F \rceil$ */

do

$p_{\text{Rand}}(z) \leftarrow \text{Dirichlet random distribution}$;

$p(z) \leftarrow p(z) + F_{\text{inc.}} n(\sigma_F) \exp[-j^2/(2\sigma_F^2)] p_{\text{Rand}}(z)$;

for n in $\{1, 2, \dots, k\}$ **do**

$p_{\text{Rand}}(z) \leftarrow \text{Dirichlet random distribution}$;

$p(z) \leftarrow p(z) + \frac{(1-F_{\text{inc.}})}{k} n(\sigma_F) \exp[-j^2/(2\sigma_F^2)] p_{\text{Rand}}(z)$;

end

end

$p(z) \leftarrow p(z) \times p_{\text{avg}}(z) / \|p(z) \times p_{\text{avg}}(z)\|_1$; /* Multiply mixture of Dirichlet random distributions by p_{avg} and normalize */

$n_{\text{trial}} \leftarrow 100$;

$\vec{n} \leftarrow \vec{0}$; /* Initialize mean for finite sample PoP */

for j in $\{1, 2, \dots, n_{\text{trial}}\}$ **do**

$\{z_1, \dots, z_M\} \leftarrow M$ samples from $p(z)$;

$N_z \leftarrow \#(z_i = z)$;

$\vec{n} \leftarrow \vec{n} + \text{Hist}(\{\frac{N_z}{M p_{\text{avg}}(z)}\}; \Delta\tilde{p}) / n_{\text{trial}}$; /* Form histogram of rescaled empirical frequencies $N_z / (M p_{\text{avg}}(z))$, with bin width $\Delta\tilde{p}$, add to mean vector \vec{n} . Each entry is the expected number of bitstrings with empirical frequency in bin. */

end

to potentially invert this finite sampling effect. In the next section we show that low-order moments of the PoPs can be estimated in an unbiased way which is robust to a finite number of samples.

7.6.1 Comparing experiment and finite-sampled model PoP

As shown in Fig. 7.6, we compare the experimental empirical PoP against PoPs obtained from analytical models, accounting for the finite sampling of the experiment. Specifically, we accrue roughly 32000 blockade-satisfying [36] bitstrings from the experimental evolution, which is $m \approx 20.6 \times$ the Hilbert-space dimension of $D = 1597$, and compare against analytical PoPs which are modified to account for this degree of finite sampling (described below). For our *ab initio* error model, we perform the full quantum simulation, then sample an equivalent number of bitstrings from the resultant state, then repeat this sampling many times to find the mean finite-sampling prediction.

For the analytical noise models, no direct quantum simulation is necessary. We first take the analytical model PoP, and sample a given normalized probability vector from it. We then sample this probability vector as many times as the experiment was sampled and reconstruct an empirical model PoP. We repeat this procedure many times, and take the average over iterations to define the numerically sampled model PoP. We emphasize that in doing so, no actual quantum simulation was needed, simply repeated resampling. The results of this resampling are in excellent agreement with analytical predictions where available (discussed below), as in Fig. 7.16.

To quantitatively compare experiment and model PoPs (as in Fig. 7.6b), we use the symmetric, standardized χ^2 -distance (CSD) defined as

$$\chi^2(\vec{x}, \vec{y}) = \frac{1}{B} \sum_i^B \frac{(x_i - y_i)^2}{(x_i + y_i)/2}, \quad (7.70)$$

where B is the number of bins in the PoP, x_i is a bin drawn from experiment, and y_i is a bin drawn from the model distribution. We calculate the CSD for each iteration of the model PoP resampling, and then define the mean CSD and its error bar as the mean and standard deviation over all such iterations.

For our comparison between the model PoPs and RUC numerics in Fig. 7.6c, we again employ the CSD as our distance, and compare our noisy RUC simulations against model PoPs. For each point in Fig. 7.6, we average over 100 RUC circuit realizations. We show additional results in Fig. 7.15, where we see the estimation correlation improves with greater

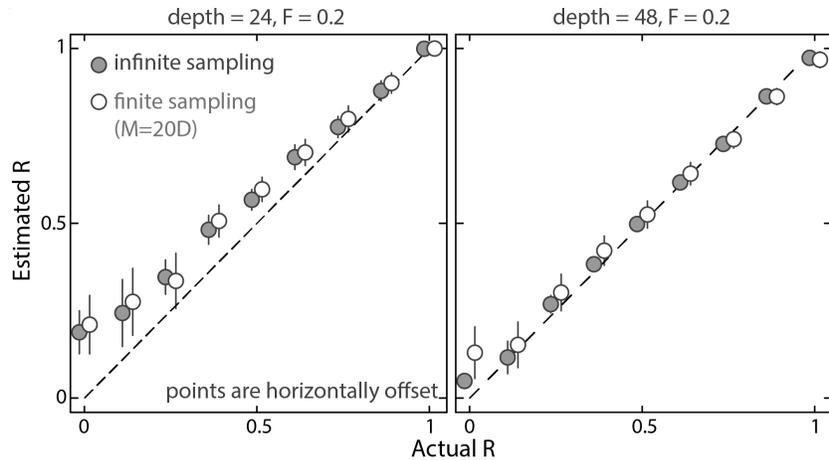


Figure 7.15: **Noise learning improves with greater depth.** We perform RUC numerics with varying fraction of global coherent error to local incoherent error, as in Fig. 7.6c, with parameters chosen such that the final fidelity is 0.2 after either depth 24 (left) or depth 48. For shorter depth circuits, the estimator correlation becomes worse, which we believe is due to a) limited scrambling, and b) not accounting for higher-order terms in the analytic noise channels (i.e. non-dilute effects of the local error channel). We also show the results for either infinite or finite sampling, which are horizontally offset from each other to better distinguish, which does not seem to have a dramatic effect on the estimation correlation.

circuit depth at an equal fidelity, which we attribute to a) limited scrambling in shallower circuits and b) not-accounting for higher-order terms in the analytic noise channels; we note that the second of these effects is solvable with model-level refinements, discussed in Secs. 7.5.4 and 7.5.4. We further find that finite sampling does not introduce significant bias in the estimator correlation, only increasing its statistical fluctuations.

In the following subsections, we describe an analytic description of finite sampling for the PoP distribution, and potential ways to undo these effects.

7.6.2 Multinomial distribution and Poisson sampling

Taking M i.i.d. samples from a probability distribution $p(z)$ yields a number n_z of occurrences of each bitstring z . The vector $\vec{n} \equiv (n_{z_1}, \dots, n_{z_D})$ of such numbers is a random variable following the *multinomial distribution* $\vec{n} \sim \text{Multinomial}(M, p(z))$ with PDF

$$P(\vec{n}; M, p(z)) = \frac{M!}{n_{z_1}! \dots n_{z_D}!} p(z_1)^{n_{z_1}} \dots p(z_D)^{n_{z_D}}, \quad (7.71)$$

where we have indexed the bitstrings from 1 to D . This is known *i.i.d. sampling* [359].

While Eq. (7.71) is an exact statistical description of the outcome of sampling, it is analytically easier to use the Poisson sampling model, which assumes that each number n_z is an

independent Poisson random variable with mean $Mp(z)$, i.e. $n_z \sim \text{Poi}(\mu = Mp(z))$ [359]. Since each number variable is independent, the constraint $\sum_z n_z = M$ is not enforced; in the multinomial distribution, there are weak anticorrelations between number variables n_z and $n_{z'}$. The Poisson sampling approximation becomes increasingly accurate with increasing sample number M and serves as a useful model for sampling [359]. Nevertheless, the full multinomial distribution is necessary for quantitatively accurate formulae in some cases which we will indicate.

7.6.3 Finite-sample estimate of PoP distribution

The Poisson sampling model will be essential to helping us understand the effects of finite sampling on reconstructing the PoP distribution. Here, we wish to estimate the distribution of probabilities $P(p)$ from empirical data. The challenge lies in the fact that the probabilities $p(z)$ are difficult to empirically estimate — the *plug-in estimate* $\hat{p}(z) \equiv n_z/M$ is a rescaled Poisson variable and is unbiased, but may exhibit large fluctuations. For simplicity of discussion, we first consider the case where $p_{\text{avg}}(z) = 1/D$ (as in RUCs) before tackling the general case. In this case, the possible values of $\hat{p}(z)$ are discrete; they are integer multiples of $1/M$, and the rescaled probability $\tilde{p}(z)$ is simply $Dp(z)$. Accordingly, we shall express our PoP as $P(Dp)$, which is system-size independent.

A standard practice to empirically estimate the PoP is to simply take the histogram of $\hat{p}(z)$. That is, one takes M samples from $p(z)$, measuring each bitstring n_z times. The empirical frequencies $\hat{p}(z) = n_z/M$ take discrete values in multiples of $1/M$ (when $p_{\text{avg}}(z)$ is nontrivial, the empirical rescaled frequencies are continuous) and estimates the PoP by the fraction of bitstrings with a given empirical frequency $\hat{P}(D\hat{p}) \equiv N_{\hat{p}}/D$.

For each bitstring z with probability $p(z)$, the empirical frequency n_z/M is a rescaled Poisson random variable. For a given empirical frequency \hat{p} , the number of bitstrings $N_{\hat{p}}$ with that empirical frequency is also a Poisson random variable with expected value

$$\hat{P}(D\hat{p} = nD/M) = \int_0^\infty d(Dp) P(Dp) p_{\text{Poi}}(n; pM), \quad (7.72)$$

where $p_{\text{Poi}}(n; pM) \equiv e^{-pM} (pM)^n / n!$ is the probability of a Poisson variable taking value n when the expected number of counts is $\mu = pM$. This prediction determines the empirical PoP with a finite number of samples. We emphasize that Eq. (7.72) is the expected value of the finite-sampling PoP, and a given experimental PoP will show statistical fluctuations about this mean. In practice, occurrences for large Dp occur with sufficiently small probability that the number of potential values may be truncated at some level. Then, as long as

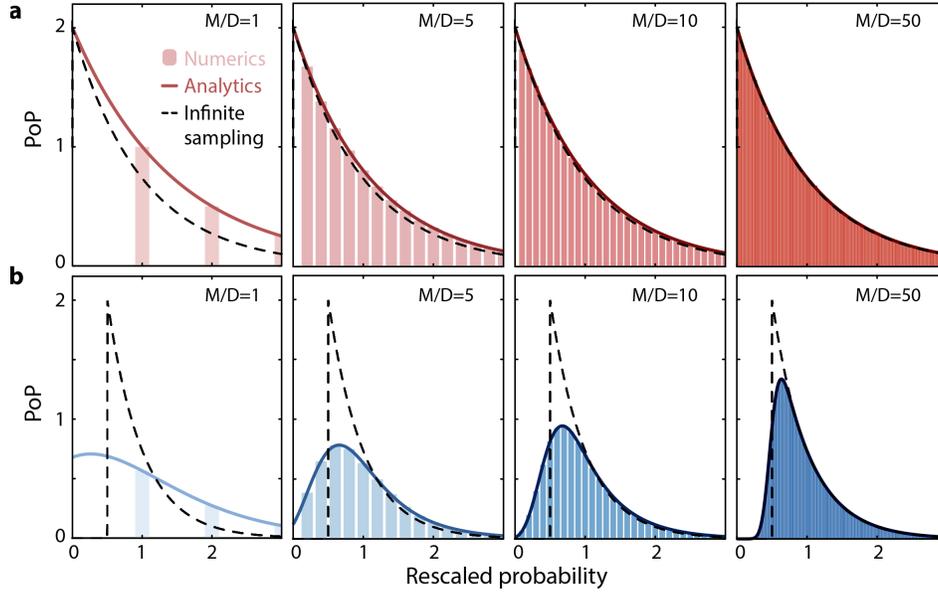


Figure 7.16: **Finite sampling effects on the probability-of-probabilities (PoP).** The PoP is nontrivially modified when reconstructed from a finite number of samples, as generally considered in Eq. (7.72). However, not all PoPs are as susceptible as others; for instance, **(A)** the Porter-Thomas distribution (Eq. (7.74)) shows relatively little sensitivity to M/D (the ratio of number of measured bitstrings divided to the total Hilbert-space dimension), while **(B)** the PoP for global depolarization noise with $F = 0.5$ (Eq. (7.75)) shows greater sensitivity due to the discontinuous PDF. In both cases, we see good agreement between numerical finite sampling and the analytical predictions Eqs. (7.74) and (7.75).

M/D is large enough, there will be enough outcomes that each PoP count is approximately independent (i.e. the normalization constraint is negligible), acting as a Poisson random variable with $\sqrt{N_{\hat{p}}}$ uncertainty.

When the factor $p_{\text{avg}}(z)$ is not constant, the empirical rescaled probability $\hat{p}/p_{\text{avg}}(z)$ takes on continuous values and we must perform histogram binning instead. Here, the expected value of the histogram is given by

$$\hat{P}(\hat{p}) \approx \sum_{n=0}^{\infty} \frac{1}{D} \sum_z \int d\tilde{p} P(\tilde{p}) p_{\text{Poi}}(n; M p_{\text{avg}}(z) \tilde{p}) \delta\left(\hat{p} - \frac{n}{M p_{\text{avg}}(z)}\right). \quad (7.73)$$

The empirical PoP $\hat{P}(\hat{p})$ is akin to (but not exactly) a convolution of the underlying PoP with the Poisson PDF. It solely depends on the underlying PoP, $P(Dp)$, as well as the ratio of samples to Hilbert-space dimension, $m \equiv M/D$. In the limit $m \rightarrow \infty$, the estimate converges to the true value $\hat{P}(D\hat{p}) \rightarrow P(Dp)$. Finite values of m show potentially substantial deviations between $\hat{P}(D\hat{p})$ and $P(Dp)$. In Fig. 7.16, we illustrate the effects

of finite sampling when the underlying PoP distribution is Porter-Thomas or arising from global depolarization errors [Eq. (7.66)]. The Porter-Thomas (PT) distribution is accurately reconstructed with modest ratios $m \gtrsim 1$. However, the global depolarization distribution shows a sharp discontinuity at $Dp = 1 - F$, which is smoothed out under finite sampling. Even large sampling ratios such as $m = 30$ show considerable deviations.

These two cases can be analytically solved. The PT distribution yields the following finite-sampling PoP

$$\hat{P}(D\hat{p})d(D\hat{p}) = \frac{1}{P_{>0}} \left(\frac{m}{1+m} \right)^{1+mD\hat{p}} d(D\hat{p}), \quad (7.74)$$

where $P_{>0} = m/(1+m)$ is the probability of observing a non-zero value of \hat{p} in Eq. (7.72), elaborated on below. Meanwhile, the global-depolarization PoP gives:

$$\hat{P}(D\hat{p})d(D\hat{p}) = \frac{1}{P_{>0}} F^{-1} e^{F^{-1}-1} \left(\frac{Fm}{1+Fm} \right)^{1+mD\hat{p}} \frac{\Gamma(1+mD\hat{p}, \frac{1-F}{F}(1+Fm))}{\Gamma(1+mD\hat{p})} d(D\hat{p}), \quad (7.75)$$

where $P_{>0} = 1 - e^{m(F-1)}$, $\Gamma(k, x) \equiv \int_x^\infty z^{k-1} \exp(-z) dz$ is the upper incomplete Gamma function and $\Gamma(k)$ is the Gamma function.

In the literature, the probability $\hat{P}(\hat{p} = 0)$ is typically not included in the empirical PoP, and we must renormalize the resulting histogram by the factor $P_{>0}^{-1}$. Indeed, if one does include $\hat{P}(\hat{p} = 0)$, the finite-sampling deviations from the PT distribution become more evident, the factor $P_{>0}^{-1}$ should not be included in Eqs. (7.74) and (7.75), and Eq. (7.72) is true without modification.

7.6.4 Inverting finite sample effects

Proof of principle: inverting with Laguerre polynomials — Remarkably, for any value of m , it is theoretically possible to invert Eq. (7.72) with Laguerre polynomials (making a technical assumption that the PoP $P(Dp)$ is square-integrable, which is true in any finite system). We expand $P(Dp)$ in terms of Laguerre polynomials:

$$P(Dp) = \sum_{k=0}^{\infty} c_k L_k(2mDp) \exp(-mDp), \quad (7.76)$$

where $L_k(x)$ is the k -th Laguerre polynomial, satisfying the orthogonality relation $\int_0^\infty dx e^{-x} L_n(x) L_m(x) = \delta_{nm}$ and c_k are the unknown coefficients to be determined. Using the above orthogonality relation, c_k can be extracted from the empirical PoP, $\hat{P}(D\hat{p}_n)$, with the formula

$$c_k = 2 \sum_{n=0}^k \binom{k}{n} (-2)^n \hat{P}(D\hat{p}_n), \quad (7.77)$$

where $\hat{p}_n = n/M$ are the discrete values of the empirical frequencies. However, this formula is not practical. It is not robust to any amount of statistical or measurement error since it weights the measured frequencies \hat{P} ($D\hat{p} = Dn/M$) with exponentially increasing and oscillatory weights in n . Nevertheless, this serves as a theoretical basis that the ideal PoP can be reconstructed from any finite-sampling PoP. For a more practical procedure, we employ a regularized least squares approach.

Inversion with regularized least squares — Our practical approach to undo the effects of finite sampling will be based on regularized least squares. We note that this method is not used for any data shown, and we instead present it for completeness and as a potential way to further improve our results.

We discretize the domain of the PoP into a discrete number of histogram bins centered at \tilde{p}_j and with width $\Delta\tilde{p}$. For concreteness, it suffices to stop at $\tilde{p} = 5$. We also discretize the finite-sampling PoP into discrete and finite histogram bins. When $p_{\text{avg}}(z) = 1/D$, its domain is naturally discrete, but this is not so for nontrivial $p_{\text{avg}}(z)$. Working in the former case for simplicity, we can rewrite Eq. (7.72) as a linear equation

$$\hat{\mathbf{P}} = \mathcal{M}\mathbf{P}, \quad (7.78)$$

$$\mathcal{M}_{n,j} = \exp(-\tilde{p}_j M) \frac{(\tilde{p}_j M)^n}{n!} (\Delta\tilde{p}). \quad (7.79)$$

When $p_{\text{avg}}(z)$ is nontrivial, one must perform histogram binning on both the range and the domain, and the matrix \mathcal{M} is instead

$$\mathcal{M}_{i,j} \approx \frac{1}{D} \sum_z \sum_{n=\lceil M(\hat{p}_i - \Delta\hat{p}/2)p_{\text{avg}}(z) \rceil}^{\lfloor M(\hat{p}_i + \Delta\hat{p}/2)p_{\text{avg}}(z) \rfloor} P(\tilde{p}_j) p_{\text{Poi}}(n; Mp_{\text{avg}}(z)\tilde{p}_j). \quad (7.80)$$

Our result Eq. (7.77) essentially gives an explicit result for \mathcal{M}^{-1} in the Laguerre polynomial basis. Unfortunately, \mathcal{M}^{-1} is extremely ill-conditioned and cannot accommodate the shot-noise in $\hat{\mathbf{P}}$. If one naively attempts to solve this linear equation, the solution is wildly oscillatory and large. In order to remedy this, we simply regularize our solution and solve the least-squares equation with a penalty term:

$$\mathbf{P}_{\text{est}} = \operatorname{argmin}_{\mathbf{y}} \left[\left(\mathcal{M}\mathbf{y} - \hat{\mathbf{P}} \right)^T A \left(\mathcal{M}\mathbf{y} - \hat{\mathbf{P}} \right) + \lambda \mathbf{y}^T B \mathbf{y} \right] \quad (7.81)$$

where we have presented this problem in its full generality and A and B are arbitrary symmetric positive-semidefinite matrices that we can suitably choose to regularize the regression. When A and $B = \hat{I}$, the first term is the well-known squared error and the

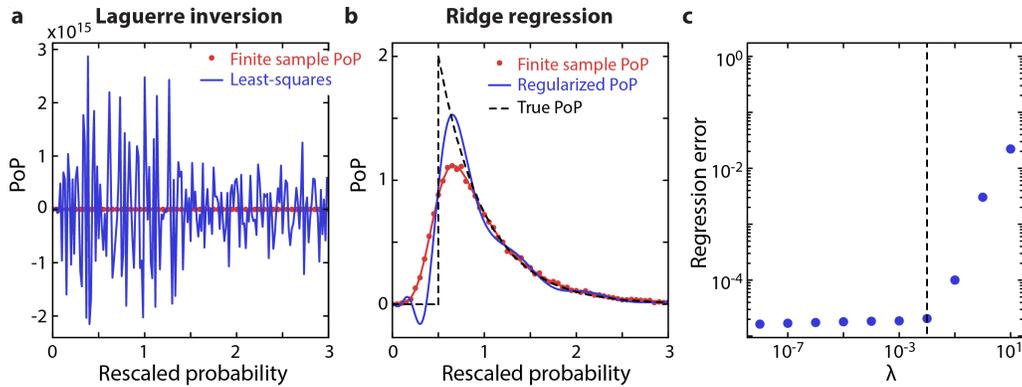


Figure 7.17: **Strategies to reverse finite-sampling effects on the PoP.** We simulate the finite-sampling effects on a global depolarization PoP, with a Hilbert-space dimension of $D = 2^{15}$, and $M = 20D$ samples. **a.** Regular least squares regression [Eq. (7.78)] is numerically unstable and leads to large, unphysical solutions (note scale of values). **b.** Regularized regression penalizes this divergence and gives an approximate reconstruction (blue) of the infinite-sampling PoP (black dashed) from the finite-sampling PoP (red points). **c.** The free parameter λ is chosen by plotting the regression error as a function of λ , displaying two regimes where the solution is respectively overfitted and overdamped. We choose λ at the transition between these regimes. Nevertheless, the reconstructed PoP is imperfect; adding further constraints, that for instance force the PoP to be always positive, may improve agreement, but we leave such avenues to future work.

second is the L_2 norm. However, we may modify A and B to enforce other constraints, such as smoothness of the solution. This is known as ridge, or Tikhonov, regression [360] and has explicit solution

$$\mathbf{P}_{\text{est}} = (\mathcal{M}^T A \mathcal{M} + \lambda B)^{-1} \mathcal{M}^T A \hat{\mathbf{P}}. \quad (7.82)$$

As an illustration, we demonstrate this in Fig. 7.17. The parameter λ is tunable, and one must choose an appropriate value. While it can be deduced analytically, here we present a user-friendly and intuitive method. In Fig. 7.17c we plot the squared error $\|\mathcal{M} \mathbf{P}_{\text{est}}(\lambda) - \hat{\mathbf{P}}\|_2^2$ as a function of the parameter λ . This shows two parameter regions. When λ is small, the squared error is nearly independent of λ . In this regime, the solution is overfitted and enforcing the constraint effectively regularizes the solution without significant effect on the squared error. When λ is large, the squared error grows quickly with λ . In this regime, the solution is overdamped and the constraint is too strong. We find that operating at the midpoint between both regimes is often optimal. Despite this, its reconstructed solutions are often imperfect, though potentially improvable through addition of additional constraints on the regression (like forcing the PoP to be always positive). Therefore, we rely on direct comparisons of the finite-sampling PoP, outlined below.

7.7 CUMULANTS OF POP DISTRIBUTIONS

As established above, due to finite-sampling effects, it can be challenging to empirically estimate the PoP. We have compared directly to finite sampled version of model PoP predictions, and in addition we have discussed potential theoretical means of reversing the finite sampling effect. Here, we discuss how we can also estimate properties of the PoP to greater precision, without reconstructing the PoP distribution itself. Specifically, the moments of the PoP can be efficiently estimated, which we may use to learn about the type of noise present in the system. Further, the moments for different error models can have different functional dependencies on the fidelity, F , meaning that analysing how moments change over multiple values of fidelity may improve predictive power.

From Section 7.5, the probability distribution can be written as $\tilde{p}(z) = \sum_i \omega_i \tilde{p}_i(z)$, where $\tilde{p}(z) \equiv p(z)/p_{\text{avg}}(z)$ is the rescaled probability and $\tilde{p}_i(z)$ are independent Porter-Thomas (i.e. Dirichlet random) distributions. Remarkably, some properties of the coefficients $\{\omega_i\}$ can be estimated from empirical samples from $p(z)$, even *without knowledge* of the constituents $\tilde{p}_i(z)$. Specifically, cumulants of the PoP are related to moments of $\{\omega_i\}$.

As above, we first discuss the case when the factors p_{avg} are constant, before generalizing. The k -th (raw, i.e. non-central) moment of the PoP $P(Dp)$ is given by

$$m_k \equiv D^{-1} \sum_z (Dp(z))^k \approx \int (Dp)^k P(Dp) d(Dp), \quad (7.83)$$

where the approximation follows by treating each value $Dp(z)$ as an independent sample from the PoP $P(Dp)$. We then consider the cumulants κ_k . For example, the first three cumulants are:

$$\kappa_1 = m_1 \quad (7.84)$$

$$\kappa_2 = m_2 - m_1^2 \quad (7.85)$$

$$\kappa_3 = m_3 - 3m_2m_1 + 2m_1^3. \quad (7.86)$$

A key property of the cumulants are that they add for sums of independent random variables [353]. That is, for $\tilde{p} = \sum_i \omega_i \tilde{p}_i$, $\kappa_k(\tilde{p}) = \sum_i \omega_i^k \kappa_k(\tilde{p}_i)$. Since each \tilde{p}_i is an independent Porter-Thomas distribution, they have identical cumulants $\kappa_k = (k-1)!$. Therefore, these k -th cumulants of the PoP are equal to the raw k -th moment of $\{\omega_i\}$

$$\frac{\kappa_k(p)}{(k-1)!} \approx \sum_i \omega_i^k. \quad (7.87)$$

As an illustration, we concretely demonstrate this with $k=2$. We have

$$\kappa_2(p) = D \sum_z p(z)^2 - 1 \approx \sum_i \omega_i^2. \quad (7.88)$$

To see this, note that $\sum_z p(z)^2 = \sum_z \omega_i \omega_j \sum_{ij} p_i(z) p_j(z)$. Since $p_i(z), p_j(z)$ are independent Porter-Thomas distributions, they are approximately orthogonal, in the sense $D \sum_z p_i(z) p_j(z) = 1 + \delta_{ij} + O(D^{-1/2})$. Using this relation, we obtain:

$$D \sum_z p(z)^2 = 1 + \sum_i \omega_i^2 + O(D^{-1/2}), \quad (7.89)$$

where we have used the fact that $\sum_i \omega_i = 1$. Remarkably, due to this approximate orthogonality, the second moment of the PoP only depends on the coefficients ω_i , up to exponentially small fluctuations.

When $p_{\text{avg}}(z)$ is nontrivial, we should replace $Dp(z)$ with $\tilde{p}(z) \equiv p(z)/p_{\text{avg}}(z)$ and the average over z as a weighted mean $\mathbb{E}_z[\cdot] = \sum_z p_{\text{avg}}(z)[\cdot]$. This modifies the moments to $m_k = \sum_z p(z)^{k-1}/p_{\text{avg}}(z)^{k-1}$.

7.7.1 Relation of hypoeponential moments to mixed state moments

The moments give key information about the mixed state. In the context of closed system dynamics, we found that the moments of the weight vector satisfy

$$\sum_i \omega_i^k \approx \text{tr}(\hat{\rho}_A^k), \quad (7.90)$$

as long as the bath B is sufficiently small.

We conjecture that the above equation also holds in the context of open system dynamics, as long as the fidelity is sufficiently high. Indeed, Eq. (7.90) is true for Haar-random measurement bases $\{|z\rangle\}$, up to subleading $O(D^{-1/2})$ corrections. If the eigenstates of $\hat{\rho}$ look “random” with respect to the computational basis $|z\rangle$, we get the desired equivalence.

The moments $\sum_i \omega_i^k$ behave differently under different noise channels. In Table 7.1, we summarize the low-order cumulants (κ_2 and κ_3) of the PoP predicted for various noise channels and their dependence on the fidelity F . By independently measuring the fidelity F (using, for example, cross-entropy benchmarking [30, 31, 162]), and the cumulants $\kappa_k(p)$, one would be able to discriminate between such noise models.

Consider first a local error channel, we have $\{\omega_i^{\text{loc.}}\} \approx \{F, (1-F)/k, (1-F)/k, \dots, (1-F)/k\}$. Therefore, its second moments

$$\sum_i (\omega_i^{\text{loc.}})^2 = F^2 + \frac{(1-F)^2}{k} \stackrel{k \rightarrow \infty}{\approx} F^2. \quad (7.91)$$

This relationship generalizes to higher moments: the largest weight F dominates, and we have $\sum_i (\omega_i^{\text{loc.}})^k \approx F^k$.

Noise channel	κ_2	κ_3
Global depolarization	F^2	$2F^3$
Local incoherent	$F^2 + (1 - F)^2/k$	$2\left(F^3 + (1 - F)^3/k^2\right)$
Global Gaussian coherent	$F/\sqrt{2}$	$2F^2/\sqrt{3}$

Table 7.1: **Low order cumulants of PoPs for various selected noise channels.**

We next consider global coherent errors. Here, the weights follow a Gaussian distribution

$$\omega_i^{\text{coh.}} = \exp(-i^2/(2\sigma_F^2))/\mathcal{Z}(\sigma_F) \quad (7.92)$$

with a width σ_F that depends on the fidelity F . When $\sigma_F \gg 1$, the normalization factor is well-approximated by $\mathcal{Z}(\sigma_F) \approx \sqrt{2\pi\sigma_F^2}$. Performing the sum, we have

$$\sum_i (\omega_i^{\text{coh.}})^2 \approx \frac{1}{4\pi\sigma_F^2} = \frac{F}{\sqrt{2}}. \quad (7.93)$$

Note the different scaling with F . In Fig. 7.18b, we verify these scaling behaviors for our two types of errors in an RUC simulation.

Note that the above prediction breaks down when $F \approx 1$. Using our Gaussian hypoexponential model above, the discrete Gaussian sums can be explicitly evaluated, giving $\sum_i (\omega_i^{\text{coh.}})^2 = \theta_3(\exp[-1/\sigma_F^2])/\theta_3(\exp[-1/(2\sigma_F^2)])^2$, where θ_3 is a Jacobi theta function. Meanwhile, using methods developed in Ref. [31], one can perform an exact calculation of the state purity in this setting, which gives $\text{tr}(\hat{\rho}^2) = (2/F^2 - 1)^{-1/2}$. As seen in Fig. 7.18c, these expressions show deviations when $F \approx 1$, with numerical data from RUC simulations more consistent with the latter expression.

In the presence of mixed noise, we further observe that since the weights factorize as $\omega_{ij}^{\text{mixed}} = \omega_i^{\text{loc.}}\omega_j^{\text{coh.}}$, so do the cumulants, in particular $\kappa_2 \approx F_{\text{loc.}}^2 F_{\text{coh.}}/\sqrt{2}$. Coupled with knowledge of the fidelity $F = F_{\text{loc.}}F_{\text{coh.}}$, one could extract the fidelities of each noise type. In Fig. 7.18B, we show proof-of-principle numerical evidence for this behaviour in a RUC.

7.7.2 From weights to Kraus operators

In addition to their relative ease of measurement, the moments of the hypoexponential weight vector, $\vec{\omega}$, can also be theoretically analyzed and related to quantities that characterize the strength of the noise channel. Given our analysis, it is natural to ask how the weight vector $\vec{\omega}$ reflects the underlying quantum channel. We find that the properties of $\vec{\omega}$ are related to the strength of the channel, but there are additional terms that we expect to be subleading. We

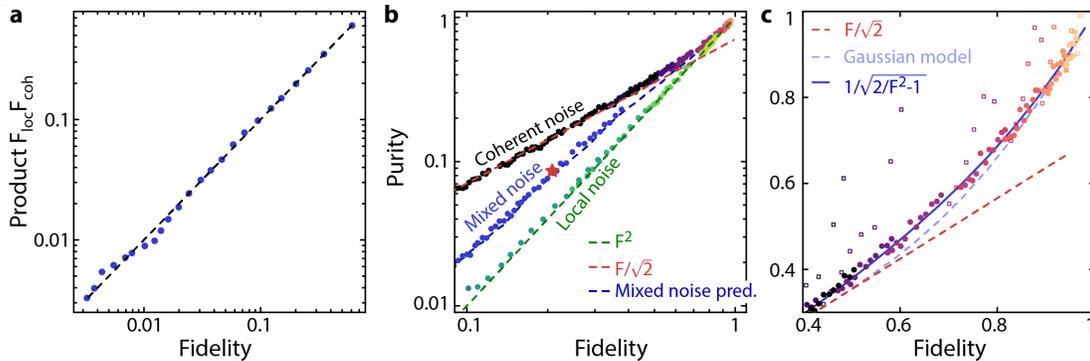


Figure 7.18: RUC simulations of mixed local and coherent errors. **a.** We first verify that errors factorize over a large range of fidelities. Here, a single instance of a random unitary circuit is presented, at a fixed proportion of local error and coherent error. As a function of circuit depth, its fidelity F decreases. We simulate the same circuit in the presence of only local or coherent errors at the same rate, obtaining the “local” and “coherent” fidelities F_{loc} and F_{coh} . The overall fidelity is well approximated by $F \approx F_{\text{loc}} F_{\text{coh}}$. This makes the concept of an error source “fraction” meaningful. **b.** Plots of the purity, or second moment, $D \sum_z p(z)^2$ for noisy RUC dynamics in the presence of global coherent or local noise, or a mixture of the two noise sources (Table 7.1). The purity due to global coherent noise asymptotically behaves as $F_{\text{coh}}/\sqrt{2}$, while the purity due to local noise behaves as F_{loc}^2 . In the presence of both types of noise, the purity behaves as $F_{\text{loc}}^2 F_{\text{coh}}/\sqrt{2}$ (blue dashed line), enabling efficient estimation of the noise content from the second moment. Our mixed error simulation is a depth-96 RUC with 12 qubits, using the same error rates as in Fig. 7.6c (precise depth marked with a red star), while our local and global coherent errors are simulations of RUCs at several error rates. Data points displayed are at depths of 24 and above. **c.** Purity from global coherent noise at high fidelities. Here, we see deviations from the asymptotic behaviour $F/\sqrt{2}$ (red dashed). In particular, it takes a value of 1 at $F = 1$. At intermediate values $0.5 < F < 0.9$, the purity deviates from the prediction from our Gaussian-hypoexponential model (blue dashed), instead agreeing with $(2/F^2 - 1)^{-1/2}$ (blue solid line), an expression obtained from a calculation of the state purity. This indicates that our model for the hypoexponential distribution [Eq. (7.67)] should be refined in this regime. In addition, to illustrate the importance of circuit depth, for the open markers we indicate the purity at depths 24 and below. These deviate from our prediction due to a lack of anticoncentration in shallow circuits.

denote the channel Φ , such that the output state is $\hat{\rho} \equiv \Phi(|\Psi_0\rangle\langle\Psi_0|) \equiv \sum_i K_i |\Psi_0\rangle\langle\Psi_0| K_i^\dagger$, where K_i are Kraus operators that describe the channel, satisfying $\sum_i K_i^\dagger K_i = \hat{I}$. The set of $\{K_i\}$ can be uniquely chosen to be orthogonal [361], in the sense that $\text{tr}(K_i K_j^\dagger) = D q_i \delta_{ij}$, such that $\sum_i q_i = 1$.

We first consider the second moment $\sum_i \omega_i^2$, which, using Eq. (7.90), approximates the purity $\text{tr}(\Phi(|\Psi_0\rangle\langle\Psi_0|)^2)$. We approximate this purity with the Haar-averaged value:

$$\mathbb{E}_{\psi \sim \text{Haar}}[\text{tr}(\Phi(|\psi\rangle\langle\psi|)^2)] = \frac{\sum_{i,j} \left(\text{tr}(K_i K_j^\dagger) \text{tr}(K_j K_i^\dagger) + \text{tr}(K_i K_i^\dagger K_j K_j^\dagger) \right)}{D(D+1)} \quad (7.94)$$

$$= \frac{D}{D+1} \left(\sum_i q_i^2 + \text{tr} \left[\Phi \left(\hat{I}/D \right)^2 \right] \right), \quad (7.95)$$

where we have used the orthogonality conditions above. The first term is the ‘‘purity’’ of the noise channel strengths p_i , while the second term is related to whether the channel is unital or not [362]. If the channel is unital, then $\Phi(\hat{I}/D) = \hat{I}/D$ so the second term has size $1/D$, which we expect to be exponentially smaller than the purity of the channel. Meanwhile, if the channel is non-unital – in an extreme case consider the amplitude damping channel such that $\Phi(\hat{I}/D) = |0\rangle\langle 0|$ – then the second term has size 1, which is larger than $\sum_i q_i^2$. In other words, the purity of the state has two contributions, given by the decoherence strength of the channel and the purity of the steady state $\Phi(\hat{I}/D)$. Assuming that $\Phi(\hat{I}/D)$ is unital or close to maximally mixed, however, we conclude that the purity of a typical input state reflects how noisy the channel is.

Higher moments give additional terms which we may not simplify, such as

$\sum_{hij} \text{tr}(K_h K_i^\dagger K_j) + \text{h.c.}$ for $k = 3$. However, we may assume that such quantities are exponentially suppressed, since the orthogonal operators K_i are obtained by the (reshaped) eigenvectors of the Choi matrix which describes the channel [361]. If these eigenvectors have sufficiently complicated bipartite structure, we expect the contributions from such terms to be exponentially suppressed. We may repeatedly apply this argument to all orders k , and find that the (expected) dominant contribution is $\sum_i q_i^k$. This also bounds the typicality error, that is the difference between the moment $\text{tr}(\Phi(|\Psi\rangle\langle\Psi|)^k)$ of a single, typical random state $|\Psi\rangle$ and the ensemble average. Indeed, we also expect our conclusion to hold as long as the initial state $|\Psi_0\rangle$ has sufficiently small overlap with the eigenstates $\{|E\rangle\}$. Therefore, we have

$$\frac{\kappa_k(p)}{(k-1)!} \approx \sum_i \omega_i^k \approx \text{tr}(\Phi(|\Psi_0\rangle\langle\Psi_0|)^k) \approx \sum_i q_i^k, \quad (7.96)$$

which therefore allows us to conclude that $\omega_i \approx q_i$, as long as the number of non-zero weights is less than the Hilbert-space dimension, i.e. $\#\{\omega_i\} \ll D$.

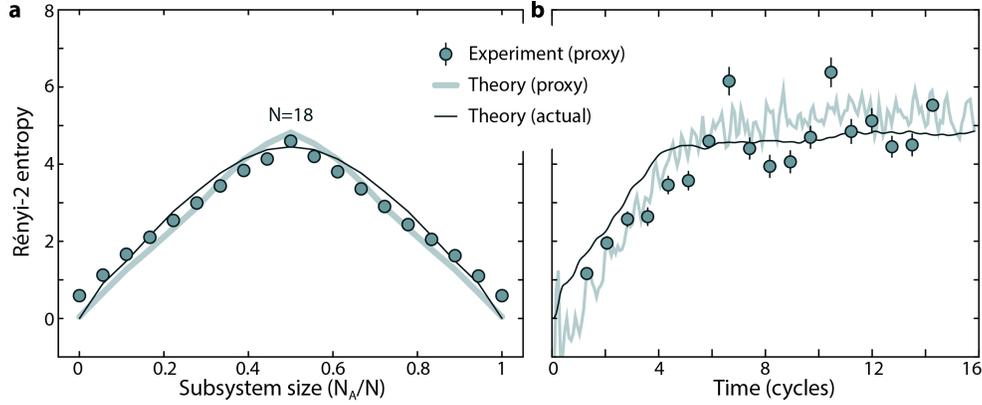


Figure 7.19: **Estimating the Rényi-2 entropy from bitstring probabilities.** **a.** Rényi-2 entropy for our experimental Rydberg array quantum simulator as a function of subsystem size following quench evolution. Proxy values are obtained from an unbiased estimator based on bitstring measurements, see Eq. (7.97). **b.** The estimator matches the growth and saturation of the true entropy as a function of time, with some underestimation at early times where bitstring probabilities are not sufficiently anticoncentrated. However, we note this representative data is slightly cherry-picked, and similarly good agreement is not seen across all times and system sizes due to effects of decoherence, limited scrambling, and arguments related to those made in Section 7.5 that the PoP for finite temperature systems is not always a hypoexponential distribution.

7.7.3 Entropy estimation from PoP moments

Our description of the PoP distribution for closed system dynamics allows a means to estimate the Rényi entanglement entropy between systems A and B . It will be convenient to consider the variance of probabilities of the bath bitstrings z_B , instead of the system bitstrings z_A . This variance estimates the Rényi-2 entropy of system A , $S^{(2)}[\rho_A] \equiv -\log_2 \text{tr}(\rho_A^2)$, through the following proxy:

$$S_{\text{proxy}}^{(2)} \equiv -\log_2 \left[\sum_{z_B} p_d(z_B) \frac{p(z_B)^2}{p_d(z_B)^2} - 1 \right] \approx S^{(2)}[\rho_A], \quad (7.97)$$

whenever $N_B > N_A$ and the measurement basis on the bath is sufficiently “unbiased” (discussed in Section 7.5). For globally pure states, $S^{(2)}[\rho_A] = S^{(2)}[\rho_B]$. However, the amount of entanglement is dictated by the size of the smaller system A , and hence we denote the entropy as $S^{(2)}[\rho_A]$. Experimental results for our Rydberg atom array quantum simulator are shown as a function of subsystem size and as a function of time for a half-chain cut in Fig. 7.19, though we note this representative dataset is slightly cherry-picked, and not all system sizes and evolution times show as good a match due to decoherence as well as more fundamental arguments described in Section 7.5.

7.7.4 Estimator of moments

Remarkably, the moments can be efficiently estimated [363] with the following estimator:

$$\hat{m}_k = D^{k-1} \sum_z \frac{(n_z)^k}{M^k} \approx D^{k-1} \sum_z p(z)^k, \quad (7.98)$$

where $n^k \equiv n(n-1)\cdots(n-k)$ is the k -th *falling power* of n and n_z is the number of occurrences of the bitstring z from M measurements. This was shown to be an unbiased estimator of m_k , with optimal sample complexity scaling. Specifically, Ref. [363] showed that the k -th Rényi entropy (obtained from m_k) can be estimated to *additive error* δ with $M = \Theta(D^{1-1/k}/\delta^2)$ samples, where D is the Hilbert-space dimension (more generally the number of possible measurement outcomes). In the same work, they showed that this scaling is optimal. Further, note that in the original reference [363], the denominator of Eq. (7.98) is M^k instead of $M^{\underline{k}}$. The former is the unbiased estimator in the Poisson sampling setting, while the latter is the unbiased estimator under realistic sampling. When $p_{\text{avg}}(z)$ is nontrivial, the moment estimator is $\hat{m}_k = \sum_z (n_z)^{\underline{k}} / (M^{\underline{k}} p_{\text{avg}}(z)^{k-1}) \approx \sum_z p(z)^{k-1} / p_{\text{avg}}(z)^{k-1}$.

We are interested in Eq. (7.98) because the low moments of the PoP can be estimated with far fewer samples (e.g. $M = O(\sqrt{D})$ for the second moment) than the entire PoP, which requires $M = O(D)$. It is illustrative to consider the simplest nontrivial case, the second moment m_2 . Our estimator $\hat{m}_2 = D \sum_z n_z(n_z - 1) / (M(M - 1))$ essentially counts the number of times a bitstring z is seen more than once. Accordingly, this estimates the *collision probability* $D \sum_z p(z)^2$, here rescaled to be typically an $O(1)$ quantity. The sample complexity is almost optimal: estimating the second moment to multiplicative accuracy δ only requires \sqrt{D}/δ^2 measurements, a manifestation of the birthday paradox. For our purposes, it is also useful to estimate its sample requirements for *additive* accuracy. Using properties of the multinomial distribution, we find that the expected uncertainty in the second moment m_2 is

$$(\Delta \hat{m}_2)^2 = \left[\frac{2M^2 - M}{M^4} D m_2 + 4 \frac{M^3}{M^4} m_3 - 2 \frac{M^2(2M - 3)}{M^4} m_2 \right] \quad (7.99)$$

$$\approx \frac{2D m_2}{M^2} + \frac{4m_3}{M} - \frac{4m_2^2}{M}, \quad (7.100)$$

where last term arises from anticorrelations between $n_z, n_{z'}$ which are not present in the Poisson sampling model. Since m_2 and m_3 are $O(1)$ quantities, this result indicates that for an additive error $|\hat{m}_2| = \delta$, the sample complexity behaves as $O(\max(\sqrt{D}/\delta, 1/\delta^2))$, where the crossover between the two scaling behaviors occurs at $M \approx D$.

To recover the result of Ref. [363] for the required number of samples to achieve a multiplicative error δ , namely $\Theta(\sqrt{D}/\delta^2)$, we focus on the second term of Eq. (7.100), then

use the inequalities $m_3 = D^2 \sum_z p(z)^3 \leq D^2 (\sum_z p(z)^2)^{3/2} = \sqrt{D} m_2^{3/2}$, and $m_2 \geq 1$ to give the desired result that $(\Delta \hat{m}_2)^2 \sim m_3/M \leq \sqrt{D} m_2^2/M$. The former inequality comes from the relation between L_p norms, $L_p \leq L_q$ for $p \geq q$, relating m_2 and m_3 to the L_2 and L_3 norms of the probability vector $p(z)$, while the latter inequality is a straightforward application of the Cauchy-Schwarz inequality. Finally, note that this result is only valid in the asymptotic regime where the second term of Eq. (7.100) dominates the first term. This is valid at a crossover value of $M \approx D$, before which the sample complexity scales as $M \sim \sqrt{D}/\delta$ instead.

Coda

We have uncovered emergent randomness arising both from partial measurement (projected ensembles) of an interacting many-body system, and from time-evolution (temporal ensembles) of that same system. Using features of the projected ensemble, we have shown a widely applicable fidelity estimation scheme which works at shorter evolution times and with reduced experimental complexity compared to existing approaches, and have demonstrated applications in quantum device comparison, Hamiltonian parameter estimation, and benchmarking the fidelity of preparing interesting quantum states. The concept of emergent randomness (projected ensembles) could provide a new framework for quantum thermalization, chaos, and complexity growth [364], which has indeed jump-started the now thriving field study of *deep thermalization* [365–373]. Such developments could enable a more flexible and standardized way of performing quantum fidelity estimation in a wide variety of quantum hardware, including trapped ions [78], superconducting qubits [276], photonic systems [374], and cold atoms and molecules in optical lattices [65], as we have subsequently theoretically demonstrated [30] with experimental showcases [31, 375].

Then, we have studied the question of a quantum system interacting with its environment through the lens of Hilbert-space ergodicity [71] both for the case of interactions between complementary subsystems, and for noisy interactions with an external environment.

In the case of closed system thermalization at infinite temperature we have analytically, experimentally, and numerically found the PoP for bitstring measurements is universally the Erlang distribution. Accordingly, there is a smooth quantum-to-classical transition with bath size from large, Porter-Thomas fluctuations to small, Gaussian fluctuations, highlighting the transition between quantum mechanical and statistical mechanical descriptions of physical systems as the number of bath configurations grow. Further, we considered systems at finite temperatures with an effective bath dimension, and showed how arbitrary moments of the PoP may be calculated. We do not yet have a closed form for the PoP at arbitrary

temperatures, but in the limit of small bath sizes we associate the finite temperature PoP with a hypoexponential distribution, with weights given by the eigenvalues of the reduced density matrix (i.e. the bipartitioned system's Schmidt coefficients).

For the case of closed system dynamics, our hypoexponential model allows us to categorize different noise models via the weight vector and produce an analytical prediction for the PoP. This enables experimentally-accessible discrimination between candidate noise models purely through their effect on the many-body bitstring PoP, without expensive numerical simulations of quantum dynamics. We have demonstrated how bespoke noise models can be found through a quantitative comparison between model predictions and experimental statistics, as has been explored in the context of random circuit sampling [376–378]. Our technique is computationally inexpensive and applies equally to analog quantum simulators and digital quantum processors. Additional quantitative tests, such as comparing the low-order moments of the PoP distributions, may enhance discrimination power while being more sample-efficient (SM). Fruitful questions remain, such as directly learning the hypoexponential weights via entanglement spectroscopy methods [379, 380]. Further, learning the weights of heavy-tailed and correlated noise distributions may guide improvements of quantum error correction schemes [342–344].

Ultimately, our results probe new types of universal emergent many-body behavior visible to high precision quantum systems. In turn, the universality we uncover has applications for both improving quantum experiments and studying fundamental quantum dynamics.

Part III

The beyond-classically-exact regime

Prelude

Quantum systems have entered a competitive regime where classical computers must make approximations to represent highly entangled quantum states [50, 162]. However, in this beyond-classically-exact regime, fidelity comparisons between quantum and classical systems have so far been limited to digital quantum devices [162, 316, 381, 382], and it remains unsolved how to estimate the actual entanglement content of experiments [72]. Here we perform fidelity benchmarking and mixed-state entanglement estimation with a 60-atom analog Rydberg quantum simulator, reaching a high entanglement entropy regime where exact classical simulation becomes impractical. Our benchmarking protocol involves extrapolation from comparisons against an *approximate* classical algorithm, introduced here, with varying entanglement limits. We then develop and demonstrate an estimator of the experimental mixed-state entanglement [72], finding our experiment is competitive with state-of-the-art digital quantum devices performing random circuit evolution [162, 316, 381, 382]. Finally, we compare the experimental fidelity against that achieved by various approximate classical algorithms, and find that only the algorithm we introduce is able to keep pace with the experiment on the classical hardware we employ. Our results enable a new paradigm for evaluating the ability of both analog and digital quantum devices to generate entanglement in the beyond-classically-exact regime, and highlight the evolving divide between quantum and classical systems.

Importantly, in general classical computers struggle to exactly represent highly entangled states [383–386], in the sense of entanglement entropy. This has raised interest in the potential of quantum devices to efficiently solve certain classically-hard problems [11, 387], but modern noisy-intermediate-scale-quantum [50, 388] (NISQ) devices are limited by experimental errors (Fig. 7.20a). This makes it a key goal to benchmark NISQ devices in the highly entangled regime where exact classical simulation becomes infeasible (Fig. 7.20b); for example, state-of-the-art classical simulation of Hamiltonian time evolution generating highly entangled states with exact global fidelity is currently limited to 38 qubits [389].

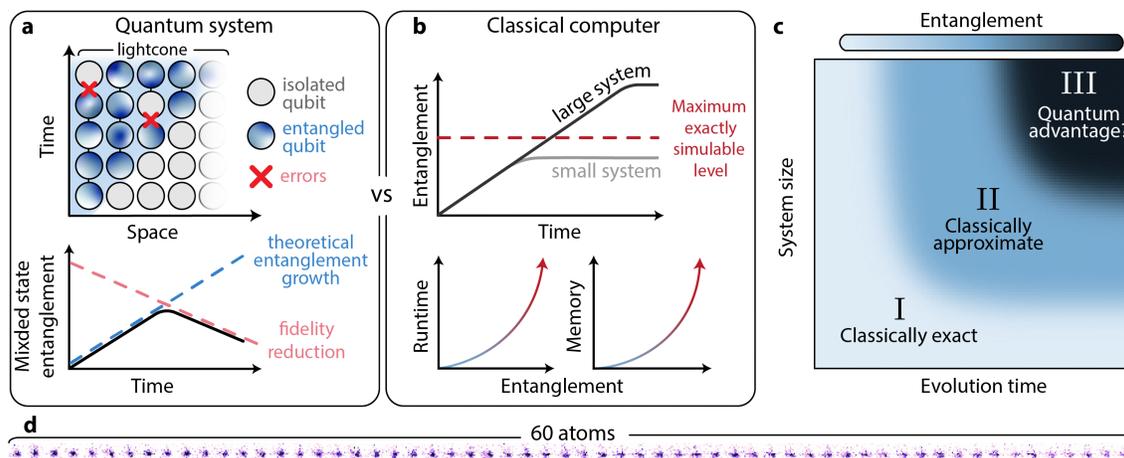


Figure 7.20: Entanglement in quantum and classical systems. **a.** In quantum systems, entanglement spreads between neighboring particles before saturating at an extensive level. However, entanglement growth is hampered by experimental errors which reduce the fidelity, limiting entanglement build-up. **b.** On the other hand, classical computers employ approximate simulation algorithms which often can only capture a limited degree of entanglement to avoid an exponential increase in cost, meaning they cannot exactly simulate dynamics at large system sizes and long evolution times. **c.** Here we compare quantum devices and classical algorithms in their ability to prepare highly entangled states using a Rydberg quantum simulator (**d**) with up to 60 atoms in a one-dimensional array (shown as a fluorescence image).

One such approach is to study the fidelity of preparing a highly entangled target state of interest [162], with several efficient fidelity estimators [29, 30, 276, 277, 390] having been introduced in recent years. However, in the beyond-classically-exact regime, these protocols have only been applied to digital quantum devices, with no such demonstrations on analog quantum simulators [391], i.e. quantum devices tailored to efficiently encode select problems of interest [36, 59–65].

In this work, we perform fidelity estimation with an analog quantum simulator targeting highly entangled states which are impractical to represent exactly on a classical computer. Our Rydberg quantum simulator [29, 36] has demonstrated [26] two-qubit entanglement fidelities of ~ 0.999 , spurring this study with up to 60 atoms [33] in a one-dimensional array (Fig. 7.20c). We stress that we target high entanglement entropy states that require an exponential number of coefficients to represent classically, as distinct from Greenberger–Horne–Zeilinger (GHZ), cluster, or stabilizer states, which are efficiently representable on a classical computer at all system sizes [392] (Fig. 7.21).

Our fidelity estimation is based on extrapolation from benchmarking against many *approximate* classical simulations, namely, matrix product state (MPS) algorithms which cap the maximum simulation entanglement to avoid the aforementioned exponential increase in

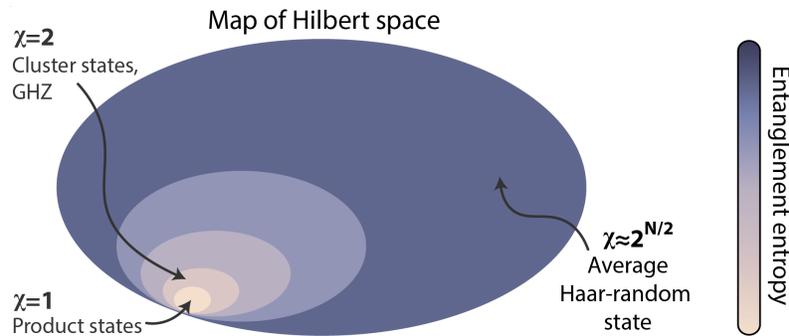


Figure 7.21: **Map of the relevant regions of Hilbert space.** For any intermediate or large system size, N , an average Haar random state in the Hilbert space will have nearly maximal entanglement entropy (up to the Page correction [324]), and correspondingly is described by an MPS with approximately maximal bond dimension, χ . By contrast, states such as GHZ state and cluster states – which are colloquially referred to as maximally-entangled – actually have very little extensive entanglement entropy, and can be written in an efficient MPS representation with only $\chi = 2$.

classical cost [392–394] (Fig. 7.20b). In one-dimension, early-time entanglement growth is system-size independent, so at short times the MPS representation is exact for essentially arbitrarily large systems. When entanglement growth surpasses the entanglement cap, the MPS is no longer a faithful reference, but we can extrapolate the fidelity through a combination of varying the system size, evolution time and simulation entanglement limit.

Using the fidelity, we derive and demonstrate a simple proxy of the experimental mixed state entanglement [72], which to date has been notoriously difficult to measure in large systems. Our proxy serves as a universal quality-factor requiring only the fidelity with, and the entanglement of, the ideal target pure state. This enables comparisons between our experiment and state-of-the-art digital quantum devices [162, 316, 381, 382, 395], with which we are competitive.

Ultimately, we compare the fidelity of our experiment against that achieved by a variety of approximate classical algorithms, including several not based on MPS. Using a single node of the Caltech central computing cluster, none of the tested algorithms is able to match the experimental fidelity in the high-entanglement regime, except for an improved algorithm we introduce, termed Lightcone-MPS. Even with this new algorithm, classical costs reach a regime requiring high-performance computing to match the experiment’s performance.

CHAPTER 8

Benchmarking large scale simulators

In Chapter 6 we introduced the concept of many-body fidelity benchmarking, predicated on our discoveries of emergent randomness in Chapter 5. Specifically, we showed that the many-body fidelity could be estimated following a short quench with two requirements: 1) a sufficient number of bitstrings were sampled from the quantum experiment, and 2) an exact classical simulation of the quantum evolution was performed to compare against. Using these tools, we demonstrated fidelity benchmarking with up to 25 atoms, close to the limit of practical exact classical simulation. Now, we ask what happens when we violate the second of these assumptions, and enter a regime beyond exact classical simulations.

In this chapter, we demonstrate how our many-body fidelity estimation techniques can be extended when system sizes grow beyond the point of classical simulability, in our case reaching as large as 60 atoms. Existing methods for benchmarking such large systems, such as patch circuits [162], explicitly fail in our setting both due to the analog control of our system, and, importantly, due to the noise sources affecting our system.

Thus we devise, develop, and deploy a new method based on performing experimental fidelity estimation with a series of *approximate* classical algorithms, each of which is erroneous, but which en masse can be used to extrapolate the correct answer. Our extrapolation procedure is based on Monte Carlo inference and ensemble learning, which we believe is generally applicable, and which should be scalable to systems of a few hundred particles. With significant cross-checks along the way, we believe these methods give accurate estimators of the actual many-body fidelity, potentially extending both the applications of Chapter 6 to the beyond-classically-exact regime, and also enabling new means in understanding the complexity of quantum simulation, as we will discuss in the remaining chapters of this thesis.

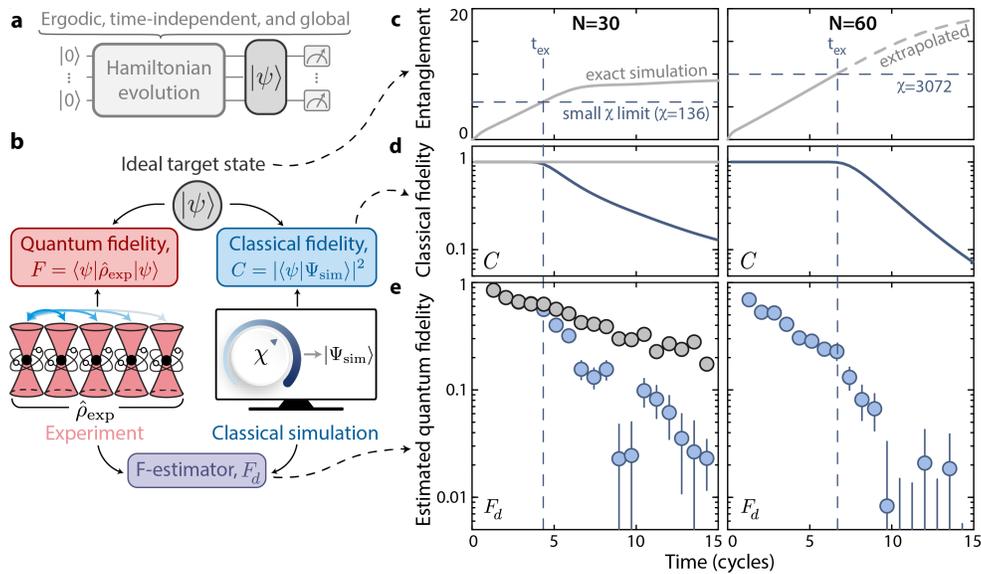


Figure 8.1: Failure of fidelity estimation with an approximate classical algorithm. **a.** We use a Rydberg quantum simulator and a classical computer to simulate a time-independent, high-temperature quench starting from the all-zero state, targeting an ideal pure state, $|\psi\rangle$. **b.** The classical algorithm is characterized by a bond dimension, χ , which limits the maximum simulable entanglement, resulting in smaller-than-unity classical simulation fidelity, C . We estimate the quantum fidelity, F , with a cross-correlation between measurement outcomes of the classical and quantum systems, termed [30] F_d . **c, d, e.** (Top) Half-cut von Neumann entanglement entropy of $|\psi\rangle$, (Middle) classical simulation fidelity, (Bottom) estimated experimental quantum fidelity. We study benchmarking against an exact simulation (gray) or an approximate simulation with limited bond dimension (blue). **c,** For a system size of $N = 30$ (left panels), using too small of a bond dimension sets a cap in the simulation entanglement. **d.** This causes the classical fidelity to fall at a time, t_{ex} , when the entanglement of the target state becomes too large. **e.** At approximately t_{ex} , the estimated experimental quantum fidelity also drops. For the largest system size, $N = 60$ (right panels), t_{ex} is well before when the entanglement saturates, even for the largest bond dimension we employ. The time-axis is normalized by the Rabi frequency (see Sec. 8.2.1).

8.1 FIDELITIES FROM APPROXIMATE ALGORITHMS

As was extensively discussed in the last Part of this thesis, a key quantity when studying quantum systems is the *fidelity* [15], $F = \langle \psi | \hat{\rho}_{\text{exp}} | \psi \rangle$, where $|\psi\rangle$ is a pure state of interest, and $\hat{\rho}_{\text{exp}}$ is the experimental mixed state. For digital devices studying deep circuits, the fidelity can be estimated via the linear-cross-entropy [162, 276], a cross-correlation between measurement outcomes of an experiment and an exact classical simulation. A modified cross-entropy, termed [30] F_d , was proposed for both analog and digital systems, and demonstrated on Rydberg [29] and superconducting [375] analog quantum simulators. F_d

is efficiently sampled (Fig. 8.23) as

$$F_d = 2 \frac{\frac{1}{M} \sum_{m=1}^M p(z_m)/p_{\text{avg}}(z_m)}{\sum_z p(z)^2/p_{\text{avg}}(z)} - 1, \quad (8.1)$$

where M is the number of measurements, z_m is the experimentally measured bitstring, $p(z)$ is the probability of measuring z with no errors following quench evolution, and $p_{\text{avg}}(z)$ is the time-averaged probability of measuring z . Importantly, $F_d \approx F$ for a wide class of physical systems, as long as the rescaled probabilities $p(z)/p_{\text{avg}}(z)$ follow the so-called Porter-Thomas distribution [30]. Still, a stringent requirement remains: access to an exact classical simulation to obtain $p(z)$, precluding direct fidelity estimation at large system sizes. We circumvent this constraint by introducing a method to estimate the fidelity by benchmarking against *approximate* classical simulations.

We consider a comparison (Fig. 8.1b) between an ideal high-entanglement target pure state, $|\psi\rangle$, the experimental mixed state, $\hat{\rho}_{\text{exp}}$, and a pure state from classical MPS simulation, $|\Psi_{\text{sim}}\rangle$. We introduce an improved MPS time-evolution algorithm using an optimal decomposition of Hamiltonian dynamics into quantum circuits [396], which we term Lightcone-MPS (see Sec. 10.2). The MPS is parameterized by a *bond dimension*, χ , that defines the maximum simulable entanglement, which scales as $\log(\chi)$. Starting from an all-zero state, we program a time-independent, global quench under the one-dimensional Ising-like Rydberg Hamiltonian (Fig. 8.1a, for Hamiltonian details see Fig. 8.4). Hamiltonian parameters lead to high-temperature thermalization [306, 307], such that describing $|\psi\rangle$ at late times requires an exponential number of classical coefficients [29].

For a system size of $N=30$ (Fig. 8.1c-e, left), we can exactly classically simulate these dynamics (Fig. 8.1d, grey); by exact, we mean the classical fidelity, $C = |\langle \Psi_{\text{sim}} | \psi \rangle|^2$, stays near unity for all times. We numerically observe the entanglement of the target state increases linearly at early times, before eventual near-saturation (Fig. 8.1c). Moreover, the estimated experimental quantum fidelity, F_d , shows apparent exponential decay due to experimental errors [29] (Fig. 8.1e, grey).

However, the situation changes when using an *approximate* classical simulation. Now, the classical fidelity begins to decay (Fig. 8.1d, blue) after the time, t_{ex} , when the ideal entanglement exceeds the limit set by the bond dimension (Fig. 8.1c, blue), meaning the classical simulation is no longer a faithful reference of the ideal dynamics. Most importantly, we find that after t_{ex} the experimental benchmarked fidelity *also deviates downwards* (Fig. 8.1e, blue), indicating that F_d no longer accurately estimates the fidelity to the ideal state. For the largest system sizes (for instance, $N = 60$ in Fig. 8.1c-e, right), t_{ex} occurs well before the entanglement is predicted to saturate, even for the largest bond

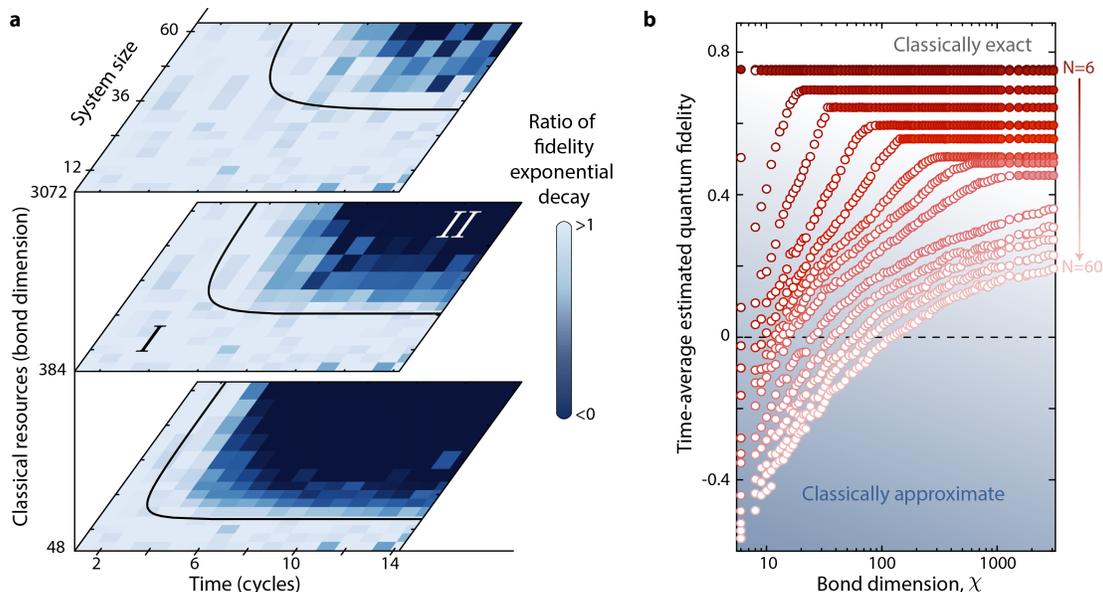


Figure 8.2: **Failure of fidelity estimation with an approximate classical algorithm.** **a.** We fit the early time fidelity decay for all system sizes and normalize the experimental fidelity measurements by these exponential values. The resulting ratio shows two clear regimes of either roughly exponential decay (white areas) or greatly superexponential decay (blue areas). These areas correspond to the regime of exact classical simulation (I) and approximate classical simulation (II) in Fig. 7.20. Notably, the boundary between regimes I and II (black line) depends on the bond dimension used in fidelity estimation. **b.** The estimated fidelity (averaged over all experimental times) increases with bond dimension (open markers), before saturating (closed markers) at a bond dimension capturing the necessary entanglement. For the largest system sizes, saturation is not achieved using the available classical resources.

dimension we can realistically use. We estimate the classical fidelity in this case using the product of MPS truncation errors [394], which we find is accurate in the regime in which we operate (Fig. 10.8).

Essentially, F_d appears to be an amalgam of both classical and quantum fidelities, only estimating the quantum fidelity to the ideal state in the limit of the classical simulation being perfect. To test this behavior for all system sizes, we study the benchmarked value of F_d averaged over all experimental times (Fig. 8.2b). Consistently we see for too small of a bond dimension (open markers), F_d is reduced. In some cases the requirement that $p(z)/p_{\text{avg}}(z)$ follow a Porter-Thomas distribution can be violated, resulting in F_d even becoming unphysically negative. As bond dimension increases, F_d rises, before reaching a *saturation bond dimension*, $\chi_0(N, t)$, which depends on system size and time (closed markers). For the largest system sizes and times, however, the saturation bond dimension is beyond the capabilities of any current classical hardware [389].

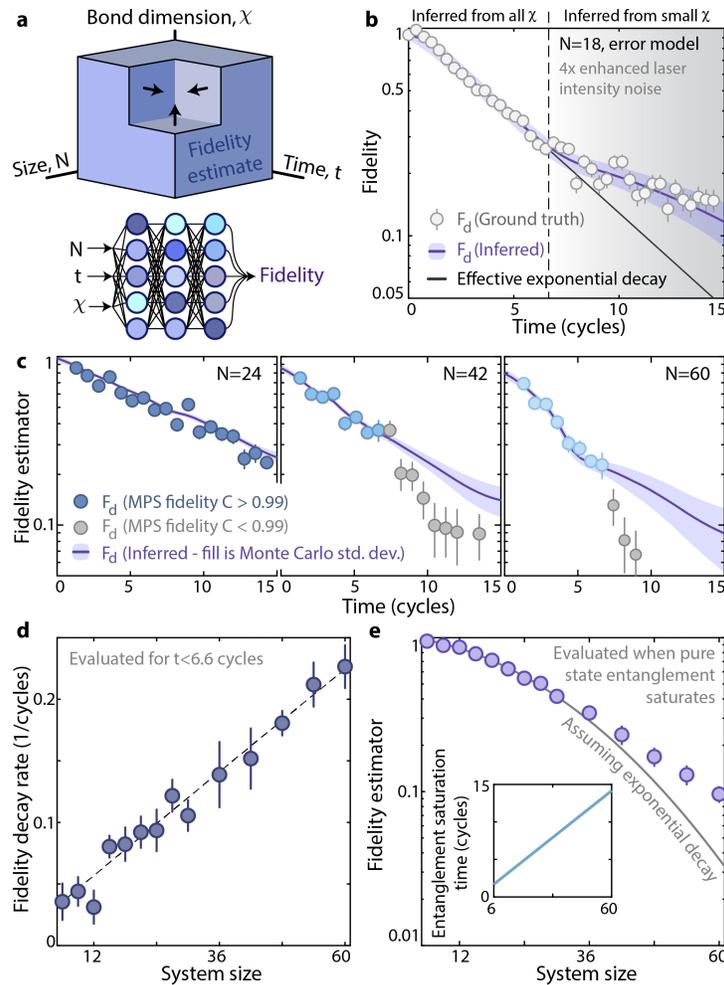


Figure 8.3: Fidelity-benchmarking a 60-atom system. **a.** We employ a Monte Carlo inference approach to extrapolate the fidelity at large system sizes and long evolution times. Specifically, we train 1500 neural networks, each instantiated with randomized (hyper)parameters, to predict F_d as a function of size, time, and bond dimension, and take the ensemble average as the predicted value. **b.** We test this procedure using error model simulations from $N = 8$ to 18 with increased laser intensity noise to emulate the fidelity expected for the experimental $N = 60$ dataset. For $t > 6.6$ cycles and $N > 15$, we only train on bond dimensions below the level necessary for exact simulation in order to mimic constraints at large system sizes. We observe two behaviors: 1) the ensemble prediction is consistent with the ground truth, and 2) the fidelity appears to follow a non-exponential form. See Sec. 8.5 for further crosschecks, as well as Sec. 8.4 for analytic evidence for the origin of the non-exponential decay due to non-Markovian noise. **c.** Experimental fidelities for N up to 60; markers are grayscale where the classical fidelity (with $\chi = 3072$) is less than 0.99. **d.** Early-time fidelity decay rate as a function of system size, consistent with linear system size scaling. **e.** Fidelity at the time (inset) where the pure state entanglement saturates, with $F_d = 0.095(11)$ at $N = 60$; the error bar is the standard error over Monte Carlo inferences added in quadrature with the underlying sampling error.

If the noise affecting the system was purely Markovian then the fidelity would decay exponentially [348] and it would be possible to measure the fidelity at early times before t_{ex} to learn the exponential decay rate, and then extrapolate in time to estimate the late-time fidelity. Indeed, we note this is an advantage of the F_d metric we employ here, because it accurately estimates the fidelity earlier than other estimators like the standard linear cross-entropy [29, 30]. However, extrapolating to late times is non-trivial in our case due to non-Markovian noise sources often affecting analog quantum systems. In particular, with analytic and numerical analysis we show that shot-to-shot Hamiltonian parameter fluctuations (e.g. laser intensity variations) induce *sub-exponential fidelity decay* at low fidelities (Theorem 1, Fig. 8.10).

Instead, we use a model-agnostic extrapolation by leveraging a large amount of data with three independent parameters: evolution time, system size, and bond dimension normalized by its saturation value (Fig. 8.3a). We can calculate F_d in seven of the octants of this parameter space - the only outlier is the high-entanglement regime of interest. We thus use a Monte Carlo inference approach by training an ensemble [397] of initially randomized neural networks to predict F_d given an input N , χ , and t ; F_d at large system sizes and long evolution times is then estimated as the ensemble average when $\chi \rightarrow \chi_0$ (Fig. 8.14). We emphasize that essentially we are simply performing curve fitting of the smoothly varying function $F_d(N, \chi, t)$, for which we can directly simulate many ground truth data.

We check that this protocol consistently reproduces fidelities at small system sizes, does not appear to overfit the experiment (Fig. 8.19), is insensitive to hyperparameters such as the neural net topology and size, and that predictions are converged as a function of the bond dimension (Fig. 8.21). We further reaffirm that our method extrapolates correctly by replicating our entire procedure in a smaller scale wherein the quantum device is replaced by numerical error model simulations up to $N = 18$ atoms (see Sec. 8.5). For $t > 6.6$ cycles and $N > 15$, the training data only consists of low bond dimensions to emulate the limitations of the large- N experimental data. Even still, the extrapolated fidelity is in excellent agreement with the ground truth data (Fig. 8.3b, Fig. 8.19), and reproduces the sub-exponential fidelity decay predicted analytically (Theorem 1).

Ultimately, we apply Monte Carlo inference to the full experimental dataset for system sizes up to $N = 60$ atoms (Fig. 8.3c; see Figs. 8.9 and 8.18 for all data). At high fidelities ($\gtrsim 0.2$), we observe nearly exponential decay, with a rate scaling linearly with system size (Fig. 8.3d). At low fidelity, however, the Monte Carlo prediction reproduces the expected sub-exponential response. We estimate the fidelity to produce the target state when the entanglement is expected to saturate (Fig. 8.3e) to be $F_d = 0.095(11)$ at $N = 60$.

To our knowledge, this is the first global fidelity estimation in the classically-inexact regime with an analog quantum simulator, and it furthermore showcases benchmarking a quantum device by extrapolating from approximate classical simulations. We expect this approach to be scalable; by studying the convergence of predicted fidelities as a function of bond dimension, our approach appears feasible for up to an order-of-magnitude more atoms than we employ here (Fig. 8.22).

8.2 EXPERIMENT DETAILS

Hamiltonian parameters are precalibrated via *ab initio* measurements, and then refined using maximum likelihood estimation (MLE) from fidelity measurements at small system sizes ($N \leq 21$), as described in Ref. [29] and Sec. 6.3.1. MLE estimated parameters are insensitive to system size, so we believe they are accurate for the larger system sizes we test here. MLE results are consistent with pre-calibrated values, and tightly optimize the target Hamiltonian used in classical simulation (Fig. 8.5a). Note that the real-space positions of atoms are taken into account when performing MLE on parameter gradients. For instance, several small islands of atoms are benchmarked over the same spatial extent as the largest 60-atom array.

The main experimental dataset was taken continuously over the course of 17 days. Parameters such as Rabi frequency, detuning, beam alignment, state preparation, and others were automatically calibrated via our home-built control architecture [29], resulting in high stability even over such a long data-taking period. The overall duty cycle of our experiment (the ratio of time spent on ‘science shots’ versus the total wall time) was roughly 36% over that period. As we load many more atoms than are needed when benchmarking small system sizes, we simultaneously excite and detect multiple non-interacting subensembles when possible, and accrue several thousand bitstrings per system size and time (Fig. 8.23a, inset).

8.2.1 Defining the experimental time unit

Throughout, we define the time unit in terms of ‘cycles,’ given by $t_{\text{cycle}} = 1/\Omega$, and thus concretely $t_{\text{cycle}} \approx 145$ ns. This timescale was chosen as Ω is the dominant energy-scale in the system (after the blockading nearest neighbor interaction energy). However, it is instructive to compare cycles versus a more natural timescale, namely the ‘unit-entanglement-time,’ t_{ebit} , required to generate 1 ebit of entanglement in the early time linear growth regime (see Fig. 8.1c). For our system parameters, we numerically find this

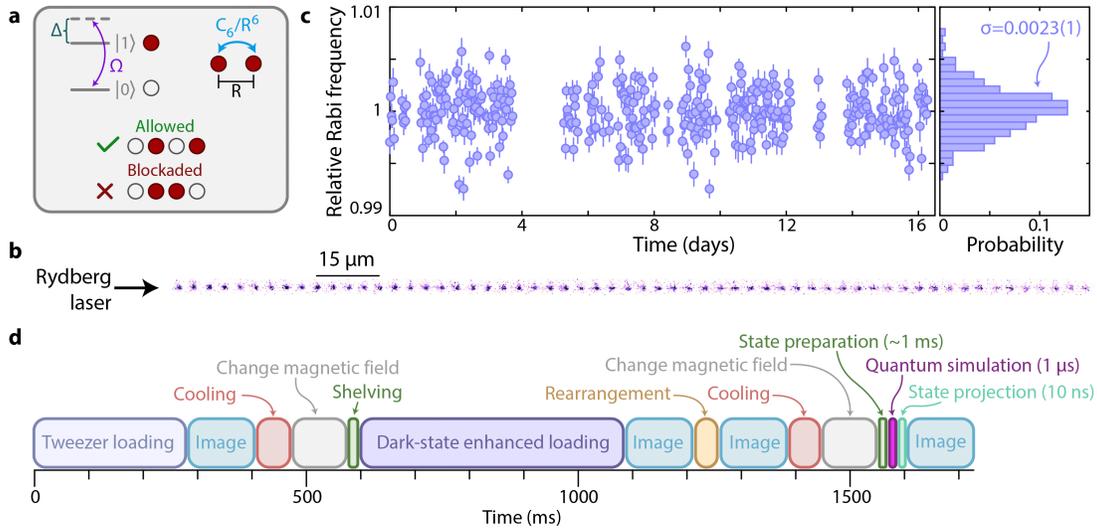


Figure 8.4: Description of the experiment. **a.** The Rydberg Hamiltonian is Ising-like, with single particle terms characterized by a detuning, Δ , and Rabi frequency, Ω , between the two qubit levels. Rydberg atoms are strongly interacting, characterized by the C_6 coefficient, and with a strength falling off approximately as $1/R^6$, where R is the interatomic distance. When two atoms are close enough together, simultaneous excitation to the Rydberg state is strongly suppressed, causing the so-called Rydberg blockade effect. We operate in a regime in which the nearest neighbor interaction strength is $\sim 13\times$ the Rabi frequency, which is strongly blockaded. **b.** We study one-dimensional chains of atoms (fluorescence image), with the Rydberg excitation laser aligned along the longitudinal array axis to minimize inhomogeneity. **c.** A classical control system [29] continuously interleaves data-taking with automated atom-based calibration of Hamiltonian parameters, resulting in highly stable operation over the course of multiple weeks. For example, the Rabi frequency is held constant with a relative standard deviation of only $0.0023(1)$. **d.** A single run of the experiment takes ~ 1.7 s, almost entirely limited by array loading and imaging. The actual time required for quantum evolution is ~ 1 μ s. We use dark-state enhanced loading [33] to reach the largest array sizes used in this work.

timescale is given by $t_{\text{ebit}} \approx 0.83 \times t_{\text{cycle}} \approx 121$ ns. This definition is convenient, because it then allows us to directly compare our analog evolution *times* against the equivalent *depth* for digital circuit evolution.

We numerically fit the early time linear entanglement growth for one-dimensional random unitary circuit (RUC) evolution, and compare against our experiment. For the gate-set of Ref. [162], we find that $t_{\text{ebit}} \approx 1.1$ layers, meaning 1 ebit of entanglement is built up in 1.1 layers. While this value is gate-set dependent, the exact numerical values are not of too much importance, but serve to demonstrate that t_{cycle} used throughout this manuscript can be treated roughly as 1 circuit depth. For further details on comparing the evolution of

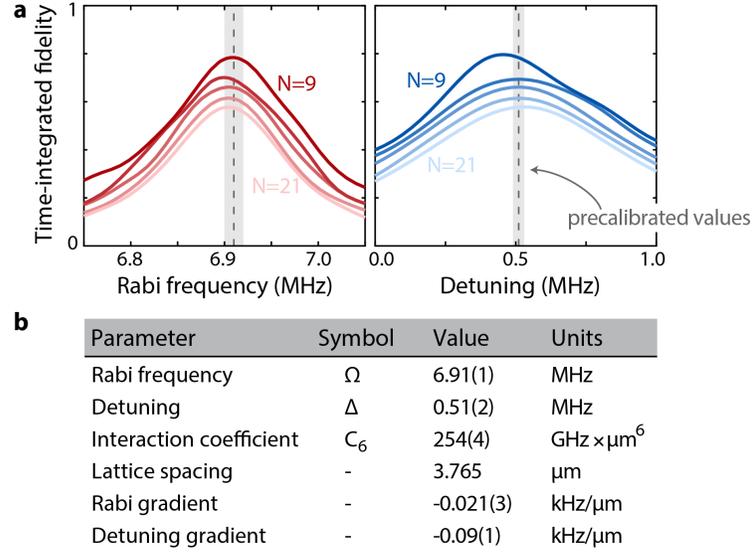


Figure 8.5: **Coherent Hamiltonian parameters.** **a.** We calibrate Hamiltonian parameters using a maximum-likelihood-estimation (MLE) technique based on varying parameters used during classical simulation, and finding the point of highest fidelity overlap with the experimental data; for further details, see Ref. [29]. Fidelities are time-integrated over all experimental measurements. MLE calibrated values are consistent for system sizes up to $N = 21$, as is expected as data taking for all system sizes (and times) is done in a randomized, interleaved fashion. Several smaller arrays are benchmarked simultaneously over the entire spatial extent of the largest 60-atom array to perform MLE on parameter gradients. **b.** Resultant MLE-calibrated Hamiltonian parameters (and gradients) used for all N in the classical simulation.

analog and digital devices, see Ref. [29].

8.3 MPS CROSS-ENTROPY BENCHMARKING

We employ a recently developed fidelity estimator [30], F_d , to evaluate the many-body fidelity of ergodic quench dynamics [70] producing large entanglement. Importantly, F_d can be efficiently sampled using only a small number of measurements as

$$F_d = 2 \frac{\frac{1}{M} \sum_{m=1}^M p(z_m)/p_{\text{avg}}(z_m)}{\sum_z p(z)^2/p_{\text{avg}}(z)} - 1, \quad (8.2)$$

where M is the number of measurements, z_m is the experimentally measured bitstring at the m^{th} repetition, $p(z)$ is the theoretically calculated probability of measuring z by a classical algorithm, and $p_{\text{avg}}(z)$ is the time-averaged probability of measuring z over a certain time window. The denominator can be efficiently calculated by importance-sampling from the MPS [398] as $\mathbb{E}_{z \sim p(z)} p(z)/p_{\text{avg}}(z)$. For small system sizes where we accrue bitstrings from

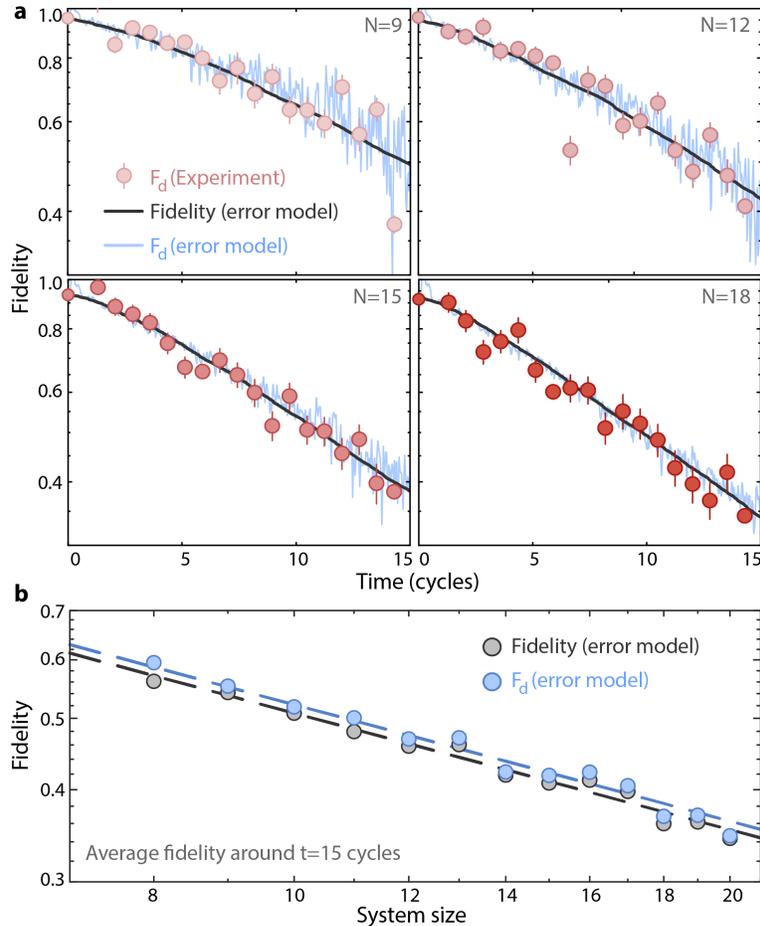


Figure 8.6: **Accuracy of the fidelity estimator.** **a.** Comparing fidelities estimated from experiment against both the true fidelity and the estimated fidelity for our error model, showing good agreement both between experiment and theory, and between the estimator F_d and the true fidelity. F_d shows systematic typicality fluctuations around the true fidelity, but these decrease in amplitude with increasing system size, see Fig. 8.23b. **b.** F_d does exhibit a systematic multiplicative offset from the true fidelity in the time regime we study due to a time-delay needed for errors to scramble and become visible [29, 30]. We find that $F_d \approx 1.02F$, regardless of system size around $t = 15$ cycles.

multiple non-interacting subarrays, we calculate F_d for each individual ‘island’ and then average their values, weighting by the number of measured bitstrings from each.

We compute $p_{\text{avg}}(z)$ via time-averaging the calculated bitstring probabilities. Specifically, for each time point t_i of the experiment, we estimate $p_{\text{avg}}(z)$ as a discrete average of the probability distributions within a window of $t_i \pm 1.4 \mu\text{s}$, with an step size of approximately 28 ns. This means around 100 points are averaged, which limits statistical sampling fluctuations. Points included in the averaging are weighted to emphasize data with similar values of classical fidelity, namely using a weight factor of $\min\left\{\frac{F_{\text{svd}}(t_i)}{F_{\text{svd}}(t_j)}, \frac{F_{\text{svd}}(t_j)}{F_{\text{svd}}(t_i)}\right\}$. Here F_{svd} is the

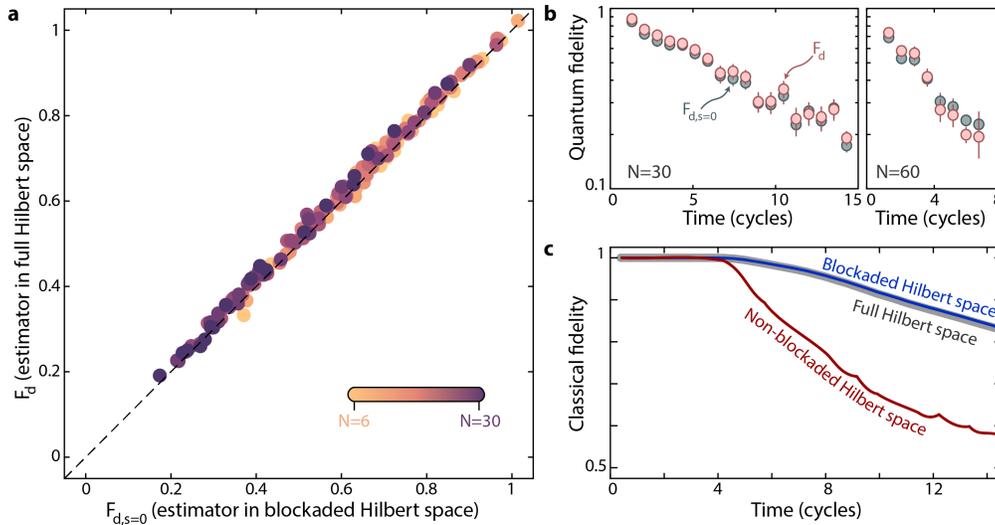


Figure 8.7: **Benchmarking in the blocked Hilbert space.** **a.** F_d versus $F_{d,s=0}$ (the fidelity benchmarked in only the blocked Hilbert space, see text) for system sizes up to $N=30$, consistent with correlation with unity slope, but with a slight ~ 0.01 offset. **b.** For exactly simulable systems, F_d typically always exceeds $F_{d,s=0}$ (left), but for larger systems (right), F_d systematically begins underestimating $F_{d,s=0}$, even before the nominal exact simulation time. **c.** We understand this behavior as arising from the MPS preferentially truncating higher blockade subspaces, leading to the non-blockaded classical fidelity dropping below unity before it is visible in the full Hilbert space fidelity. This implies $F_{d,s=0}$ stays a more accurate estimator slightly longer than F_d . **d.** Probability to be in the blocked subspace for both experiment and simulation. Note the different color-scale axes for the simulation and experiment, as many experimental errors can lead to apparently non-blockaded measurements.

product of MPS truncation errors [392, 394], which we use to approximate the classical accuracy, and defined as

$$F_{\text{svd}} = \prod_i \sum_{\alpha=1}^{\chi} s_{i,\alpha}^2, \quad (8.3)$$

where i runs over all MPS steps involving Schmidt value truncations, and $s_{i,\alpha}$ are the Schmidt values at truncation step i . Here, we assume the wave function is normalized where $\sum_{\alpha=1}^{\infty} s_{i,\alpha}^2 = 1$. While this estimation is only approximate, it can be extremely accurate when successive truncations are independent [394]. We describe this approximation in more detail in Section 10.3 and Fig. 10.8. We weight in this way under the hypothesis that averaging classical simulations of very different accuracies (as typified by F_{svd}) will lead to a poorer estimation of $p_{\text{avg}}(z)$.

To account for the Rydberg blockade mechanism, we have the option to slightly modify [29] the F_d formula. Indexing the different blockade sectors by the number of blockade violations

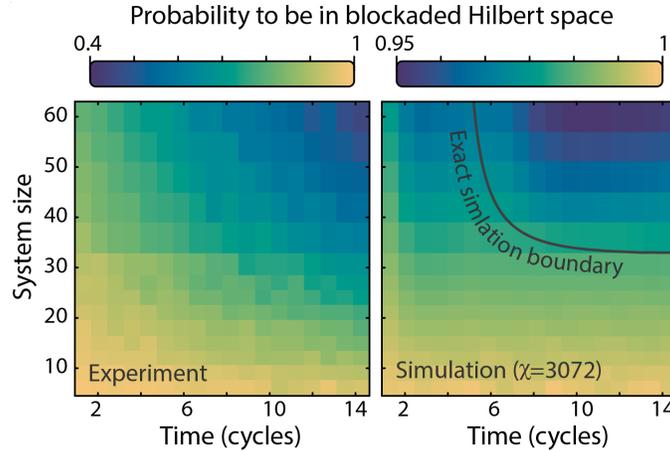


Figure 8.8: **Experimental and numerical blockade probabilities.** Probability to be in the blockaded subspace for both experiment and simulation. Note the different color-scale axes for the simulation and experiment, as many experimental errors can lead to apparently non-blockaded measurements.

in that sector, s , we can write an approximator for F_d to get an estimator for the blockade sector, $F_{d,s=0}$, only:

$$F_{d,s=0} = B_{\text{thy}} B_{\text{exp}} \left(2 \frac{\frac{1}{M'} \sum_{m=1}^{M'} p'(z'_m) / p'_{\text{avg}}(z'_m)}{\sum_z p'(z)^2 / p'_{\text{avg}}(z)} - 1 \right). \quad (8.4)$$

Here B_{exp} (B_{thy}) is the total probability for an experimental (simulation) bitstring to be in the blockaded Hilbert space, $s = 0$, which further redefines the normalized probabilities, $p'(z) = p(z) / B_{\text{thy}}$. Note here M' is the number of bitstrings, z'_m , measured in the blockaded Hilbert space.

Importantly, we find $F_{d,s=0}$ is more robust against the failure of the classical simulation algorithm, as compared to $F_{d,s>0}$. This is because MPS truncations more strongly affect the $s > 0$ sectors in ways which are not necessarily visible only by looking at the MPS truncation fidelity (Fig. 8.7bc).

For this reason, when estimating the quantum fidelity, we employ a two-pronged approach. Where $F_{\text{svd}} > 0.99$ we approximate F_d as $F_{d,s=0}$, which from measurements at small system sizes we find is always conservative (Fig. 8.7a). Utilizing $F_{d,s=0}$ allows us to more confidently directly measure the fidelity out to intermediate times at the largest system sizes, and measure the according effective decay rate (Fig. 8.3c). However, where $F_{\text{svd}} < 0.99$, we estimate F_d from our Monte Carlo inference procedure (described below), which is trained directly on the values of F_d .

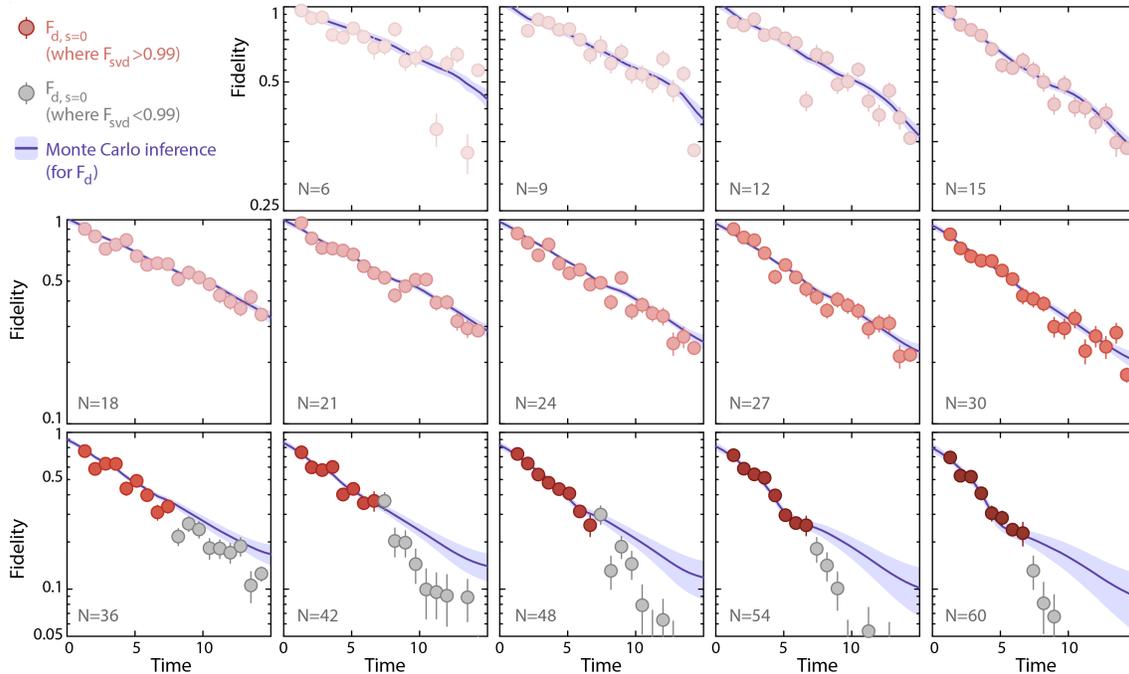


Figure 8.9: **Fidelity benchmarking for all system sizes and times.** Fidelity (quantified by $F_{d,s=0}$ with $\chi = 3072$) as well as Monte Carlo inference for all system sizes. Grayscale markers are those for which the classical fidelity is less than 0.99.

To be concrete all data shown as markers in Figs. 8.1 and 8.3 is $F_{d,s=0}$ for $\chi = 3072$. We summarize all of the data, both $F_{d,s=0}$ and the Monte Carlo fidelities, in Fig. 8.9.

8.4 NON-EXPONENTIAL FIDELITY DECAY

In Fig. 8.3, we experimentally observe that the many-body fidelity appears to decrease exponentially at intermediate times, before bending sub-exponentially at late times, a behavior which seems to grow stronger with increasing system size (also visible in Fig. 8.9). We believe this behavior originates from globally-correlated, non-Markovian Hamiltonian parameter fluctuations, which we note are often found in analog quantum simulators. In particular, a dominant noise source in our experiment is shot-to-shot variation of the Rabi frequency, for instance arising from fluctuations of the driving laser intensity, which we measure to be around $\sim 0.5\%$, corresponding to a shot-to-shot variation of $\sigma \sim 40$ kHz (Fig. 8.10a). Here, we present analytical and numerical support that such an error source can induce the observed fidelity decay behavior.

8.4.1 Fidelity response to coherent Hamiltonian errors

We discover a surprising dependence of the fidelity between states subject to coherent Hamiltonian errors.

Theorem 1 (Gaussian fidelity response to coherent errors). *Let $|\Psi(t, \theta)\rangle \equiv \exp(-i(\hat{H}_0 + \theta\hat{V})t)|\Psi_0\rangle$ denote the initial state $|\Psi_0\rangle$ quench-evolved by a translationally-invariant local Hamiltonian $\hat{H}_0 + \theta\hat{V}$ for a time t , where the perturbation $\hat{V} = \sum_x \hat{V}^x$ is a sum of local terms and θ quantifies the perturbation strength. We assume the Eigenstate Thermalization Hypothesis (ETH), i.e. that the expectation values $\langle E|\hat{V}^x|E\rangle$ of each term \hat{V}^x with respect to the eigenstates $|E\rangle$ of \hat{H}_0 are independent random samples from a fixed distribution. Then in the limit of large system size N and long time t , the fidelity F between two such states $|\Psi(t, \theta_1)\rangle$ and $|\Psi(t, \theta_2)\rangle$ is a Gaussian function of $(\theta_1 - \theta_2)t$:*

$$F(\theta_1, \theta_2) \equiv |\langle \Psi(t, \theta_1) | \Psi(t, \theta_2) \rangle|^2 = \exp(-N\lambda(\theta_1 - \theta_2)^2 t^2 / 2), \quad (8.5)$$

for some constant λ defined in Eq. (8.14).

This result is illustrated in Fig. 8.10b and will be instrumental to showing our observed non-exponential fidelity decay (Corollary 1). This fidelity response is related to the Loschmidt echo, and our result is consistent with existing numerical observations in e.g. Refs. [399–401].

Proof. Without loss of generality, it suffices to let $\theta_1 = 0$ and $\theta_2 = \theta$. We compute the Taylor series of $F(\theta) \equiv F(0, \theta)$ and show that for large N , it agrees to all orders with the Gaussian series

$$\exp\left(-\frac{\theta^2 t^2}{2\sigma^2}\right) = 1 - \frac{1}{2}\left(\frac{\theta t}{\sigma}\right)^2 + \frac{1}{8}\left(\frac{\theta t}{\sigma}\right)^4 + \dots \quad (8.6)$$

It is instructive to compute the first non-trivial term: the second derivative

$$\partial_\theta^2 F(\theta)|_{\theta=0} = \left[\partial_\theta^2 \langle \Psi(t) | \Psi(t, \theta) \rangle \langle \Psi(t, \theta) | \Psi(t) \rangle \right. \quad (8.7)$$

$$+ 2(\partial_\theta \langle \Psi(t) | \Psi(t, \theta) \rangle)(\partial_\theta \langle \Psi(t, \theta) | \Psi(t) \rangle) \quad (8.8)$$

$$\left. + \langle \Psi(t) | \Psi(t, \theta) \rangle \langle \partial_\theta^2 \langle \Psi(t, \theta) | \Psi(t) \rangle \rangle \right] \Big|_{\theta=0}$$

In order to evaluate these derivatives, we use the relation [402]:

$$\begin{aligned} \partial_\theta \exp(i(\hat{H}_0 + \theta \hat{V})t)|_{\theta=0} &= \int_0^t d\tau \exp(i\hat{H}_0\tau)(i\hat{V}) \exp(i\hat{H}_0(t-\tau)) \\ &= it \sum_E V_{EE} |E\rangle \langle E| e^{iEt} \\ &\quad + i \sum_{E \neq E'} V_{EE'} |E\rangle \langle E'| \frac{e^{iE't} - e^{iEt}}{i(E' - E)}, \end{aligned} \quad (8.9)$$

where $|E\rangle$ are the eigenstates of \hat{H}_0 with eigenvalues E , and we have defined $V_{EE'} \equiv \langle E' | \hat{V} | E \rangle$. For our purposes, Eq. (8.9) is valid even when \hat{H}_0 has degeneracies: the initial state $|\Psi_0\rangle$ projects any degenerate subspace onto a single eigenstate $|E\rangle$.

The first term in Eq. (8.9) grows linearly with time t , while the summands in the second term oscillate and are sub-dominant. We subsequently neglect such terms. We obtain a similar expression for the second derivative:

$$\partial_\theta^2 \exp(i(\hat{H}_0 + \theta \hat{V})t) \approx -t^2 \sum_E V_{EE}^2 |E\rangle \langle E| e^{iEt}. \quad (8.10)$$

This gives:

$$\begin{aligned} \partial_\theta^2 F(\theta)|_{\theta=0} &\approx -2t^2 \left[\sum_E |c_E|^2 V_{EE}^2 - \left(\sum_E |c_E|^2 V_{EE} \right)^2 \right] \\ &\equiv -2t^2 (\langle \hat{V}^2 \rangle - \langle \hat{V} \rangle^2) \\ &\equiv -2t^2 \kappa_2(\hat{V}), \end{aligned} \quad (8.11)$$

where $|c_E| \equiv |\langle E | \Psi_0 \rangle|^2$ are the populations of $|\Psi_0\rangle$ in the energy eigenbasis. We interpret the quantities $\sum_E |c_E|^2 V_{EE}$ and $\sum_E |c_E|^2 V_{EE}^2$ as averages of the vector V_{EE} and V_{EE}^2 with respect to the distribution $|c_E|^2$: we denote them $\langle \hat{V} \rangle$ and $\langle \hat{V}^2 \rangle$ respectively; the expression $\langle \hat{V}^2 \rangle - \langle \hat{V} \rangle^2$ is the variance, which we denote κ_2 for reasons made clear below.

Within the approximation made in Eq. (8.9), $\partial_\theta^n F(\theta)|_{\theta=0} = 0$ for all odd n . The next term can be expressed in terms of the moments $\langle \hat{V}^k \rangle$ as

$$\begin{aligned} \partial_\theta^4 F(\theta)|_{\theta=0} &\approx 2t^4 (\langle \hat{V}^4 \rangle - 4\langle \hat{V}^3 \rangle \langle \hat{V} \rangle + 3\langle \hat{V}^2 \rangle^2) \\ &= 2t^4 [\kappa_4(\hat{V}) + 6\kappa_2(\hat{V})^2], \end{aligned} \quad (8.12)$$

where $\kappa_4(\hat{V}) \equiv \langle (\hat{V} - \langle \hat{V} \rangle)^4 \rangle - 3\langle (\hat{V} - \langle \hat{V} \rangle)^2 \rangle^2$ is the *fourth cumulant* of \hat{V} . Crucially, $\hat{V} = \sum_x \hat{V}^x$ is the sum of identical local terms: the Eigenstate Thermalization

Hypothesis (ETH) predicts that their eigenstate overlaps V_{EE}^x fluctuate about a smooth function of energy [328]. Furthermore, for sufficiently high-temperature states, correlations $\langle \hat{V}^x \hat{V}^y \rangle$ are short-ranged, and we can treat each \hat{V}^x as an effectively independent random variable.

We utilize a further property of cumulants: they are additive, obeying $\kappa_n(\sum_x \hat{V}^x) = \sum_x \kappa_n(\hat{V}^x)$. Therefore, all cumulants scale linearly with N . When N is large, $\kappa_4(\hat{V}) \ll \kappa_2(\hat{V})^2$ and we can neglect the κ_4 term in Eq. (8.12). This gives the expression

$$F(\theta) \approx 1 - \frac{1}{2} 2\kappa_2(\hat{V})(\theta t)^2 + \frac{1}{8} (2\kappa_2(\hat{V}))^2 (\theta t)^4 + \dots, \quad (8.13)$$

which agrees with Eq. (8.6) when we identify

$$\sigma^{-1} = 2\kappa_2(\hat{V}) = N\kappa_2(\hat{V}^x) \equiv N\lambda. \quad (8.14)$$

In fact, all higher cumulants are subleading to the κ_2 contribution. As long as they can be neglected, the series Eq. (8.6) and Eq. (8.13) agree *to all orders*. For example, the $(\theta t)^6$ term of (8.13) is $-2(\kappa_6(\hat{V}) + 30\kappa_4(\hat{V})\kappa_2(\hat{V}) + 60\kappa_2(\hat{V}))^2/6! \approx -(2\kappa_2(\hat{V}))^3/48$, in agreement with the sixth-order term of Eq. (8.6).

We prove this equivalence to all orders using properties of *Bell polynomials*, B_k (where k is a polynomial index) [403]. The n -th derivative of the fidelity can be expressed as

$$\partial_\theta^n F|_{\theta=0} = (-1)^{n/2} (\theta t)^n \sum_{k=0}^n (-1)^k \binom{n}{k} \langle \hat{V}^{n-k} \rangle \langle \hat{V}^k \rangle. \quad (8.15)$$

We can rewrite this in terms of the cumulants $\kappa_1, \dots, \kappa_k$ with the cumulant-moment relation [404]

$$\langle \hat{V}^k \rangle = B_k(\kappa_1, \kappa_2, \dots, \kappa_k). \quad (8.16)$$

While the full expression is challenging, we are only interested in the coefficient of the $\kappa_2(\hat{V})^{n/2}$ term in Eq. (8.15). We first note that $\langle \hat{V}^k \rangle$ contains a $\kappa_2(\hat{V})^{k/2}$ term only for even k . To determine its coefficient, we use the recursion relation [403] $\partial_{\kappa_2} B_k = \binom{k}{2} B_{k-2}$ to conclude that $\langle \hat{V}^k \rangle = \dots + (k!/k!!)\kappa_2(\hat{V})^{k/2} + \dots$, where $k!! = k(k-2)\dots 2$ is the *double factorial* of k . Returning to Eq. (8.15), we conclude that $\partial_\theta^n F|_{\theta=0} = \dots + a_n \kappa_2(\hat{V})^{n/2} (\theta t)^n + \dots$, where

$$a_n = (-1)^{n/2} \sum_{k \text{ even}} \binom{n}{k} \frac{k!(n-k)!}{k!!(n-k)!} \kappa_2^{n/2}$$

$$= (-2)^{n/2} \frac{n!}{n!!} \kappa_2^{n/2}, \quad (8.17)$$

which equals the n -th derivative $\partial_\theta^n \exp(-\theta^2/(2\sigma^2))|_{\theta=0} = (-1)^{n/2} \frac{n!}{n!!} \sigma^{-n}$, thus concluding our derivation of the Gaussian fidelity response [Eq. (8.5)].

8.4.2 Non-exponential decay for Rydberg quantum simulators

The non-exponential fidelity decay immediately follows from Theorem 1.

Corollary 1. *When the parameter value θ follows a Gaussian probability distribution $P(\theta)$ with mean θ' and standard deviation σ*

$$P(\theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\theta - \theta')^2}{2\sigma^2}\right), \quad (8.18)$$

Theorem 1 and a straightforward integration give a universal form for the fidelity of the target state $|\Psi(t, \theta_0)\rangle$ to the mixed state $\rho(t) \equiv \int d\theta P(\theta) |\Psi(t, \theta)\rangle \langle \Psi(t, \theta)|$.

$$F = \int d\theta P(\theta) F(\theta_0, \theta) = \frac{1}{\sqrt{1 + \Lambda}} \exp\left(-\frac{\delta^2}{2} \frac{\Lambda}{1 + \Lambda}\right), \quad (8.19)$$

where $\Lambda \equiv Nt^2\lambda\sigma^2$ and $\delta = (\theta' - \theta_0)/\sigma$ is the overall relative parameter miscalibration.

This functional form leads to markedly different behavior as compared to simple exponential decay, as it becomes essentially like a power-law at late times (Fig. 8.11d). In simulation, we benchmark an $N = 16$ atom system experiencing only global intensity fluctuations – but otherwise no noise – and see an excellent agreement between the simulated fidelity, and a fit to Eq. (8.19) with free λ . Notably, we see Eq. (8.19) exhibits three distinct regimes: super-exponential behavior at very early time, effectively simple-exponential behavior at intermediate times, and finally sub-exponential behavior at late times. Notably, by studying multiple system sizes, we find the divergence between Eq. (8.19) and an effective exponential decay model is mainly a function of fidelity, not time, and becomes visible when the fidelity reaches $F \sim 0.2$, which in our dataset is only the case for the highest N (Fig. 8.3).

As a quantitative comparison, we define an effective model of the fidelity as

$$F \approx F_0 \exp(-\gamma t) \frac{1}{\sqrt{1 + \lambda' t^2}} \exp\left(-\frac{\delta^2}{2} \frac{\lambda' t^2}{1 + \lambda' t^2}\right), \quad (8.20)$$

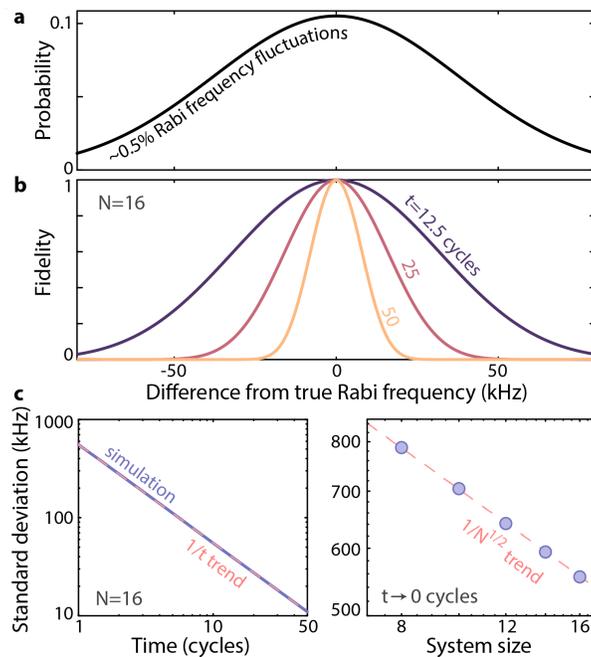


Figure 8.10: **Origins of the non-exponential fidelity decay.** **a.** A major noise source for our platform is Rabi frequency fluctuations (for instance arising from variation in the driving laser intensity), which have a shot-to-shot standard deviation of $\sim 0.5\%$ (~ 40 kHz). **b.** For global, correlated intensity noise, the fidelity is Gaussian with respect to error in the Rabi frequency. **c.** The standard deviation of this Gaussian dependence decreases as $1/t\sqrt{N}$.

that is, Eq. (8.19) multiplied by an exponential decay factor to account for the other noise sources, where we have redefined the free parameter $\lambda' \equiv \lambda N \sigma^2$. We perform error model simulations from $7 < N < 21$ for two cases: 1) for only global, shot-to-shot Rabi and detuning noise, 2) for all other noise sources. Then, assuming multiplicativity of the error sources (Fig. 8.12a), from case (1) we fit the expected behavior of λ' , and from case (2) we fit the expected behavior of γ . We extrapolate the observed parameters all the way to $N = 60$, and compare against the experiment (inset of Fig. 8.11d). Even assuming perfect calibration ($\delta = 0$) we see general agreement between the fidelity decay curve of the error model prediction and the experiment.

However, the relative proportion of exponential decay versus non-exponential decay is difficult to calibrate versus experimental data at small system sizes, as this requires evolving out to very late times when other error sources (such as finite recapture probability of atoms in their traps) become noticeable. In addition, we expect that Hamiltonian parameter calibration is likely imperfect, at the very least up to our uncertainty in the Hamiltonian

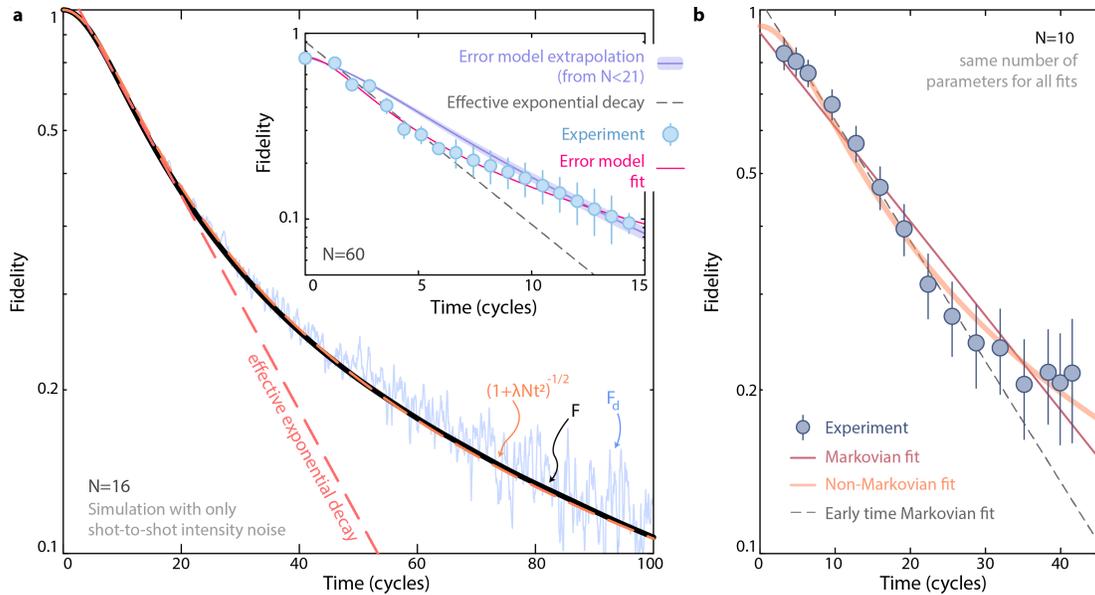


Figure 8.11: Evidence of non-exponential decay. a. The behavior described in Fig. 8.10 leads to an analytic solution for the fidelity decay of $F \sim (1 + \lambda N t^2)^{-1/2}$, for free parameter λ . This decay is characterized by super-exponential curvature at early times, and sub-exponential curvature at late times. In the inset, we use our error model with all error sources enabled, and extrapolate to the expected fidelity behavior at $N = 60$, seeing generally good agreement with the experiment, which is improved by directly fitting and allowing unity-level changes in the extrapolated error model parameters. **b.** Besides the data shown for large system sizes in Fig. 8.3, we show data for $N = 10$ evolved out to late times and low fidelities, which we see is clearly better described by the non-Markovian decay formula over an exponential.

parameters. As a general sanity check then, we fit with γ , λ' , and δ free. Fitting $N = 60$ (with F_0 fixed by the separately calibrated single-atom preparation fidelity), we find $\delta \approx 0.2$, $\gamma \approx 1/4 \times \gamma_0$ and $\lambda' \approx 4 \times \lambda'_0$, where γ_0 and λ'_0 are the extrapolated predictions from error model. We note $\delta \approx 0.2$ corresponds to a potential error of ≈ 10 kHz in the calibrated Rabi frequency, which is within our measured error bar (Fig. 8.5b).

We also note that other error sources, in particular atomic position fluctuations, also likely have some contribution to the non-exponential decay behavior, which may improve the agreement above. However, the analytical expression for their effect becomes quite involved, and so we disregard them here for simplicity, and leave a more in-depth analytic and numerical study of the various error sources to future work.

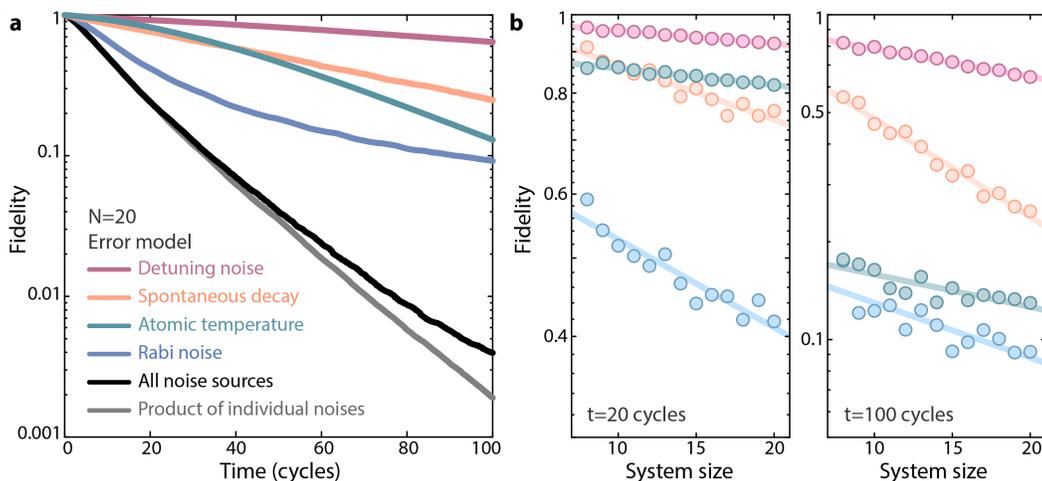


Figure 8.12: **Impact of different error sources on fidelity.** **a.** Simulated fidelity decay for a variety of choices of noise sources affecting our Rydberg quantum simulator. **b.** Fidelities for error sources in **a** at short (left) and long (right) times. Solid lines are exponential fits as a function of system size with free offset at zero time to account for non-exponential decay effects. Clear non-exponential signatures are seen for multiple error sources.

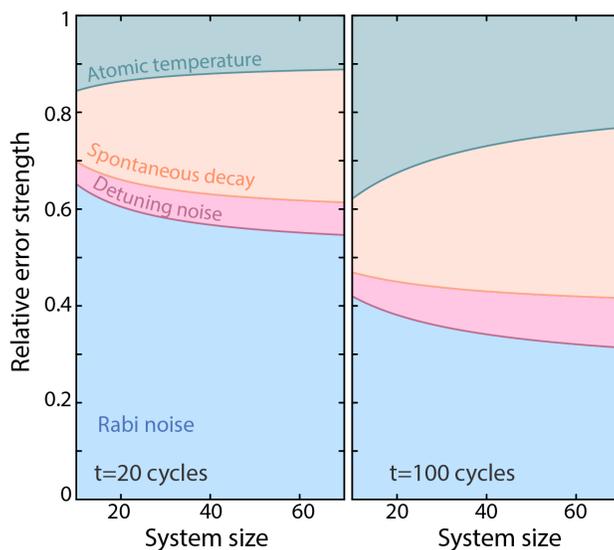


Figure 8.13: **Relative error strengths.** We use the fits in Fig. 8.12b to estimate the relative contribution of different error sources for increasing system size. Here relative error contribution is defined as $\log(F_i)/\log(F)$, where the subscript, i , indexes the error source in question, and $F = \prod_i F_i$, the product of all individual errors. Note that collective Rydberg decay processes are not accounted for here, and may further lower the fidelity measured for spontaneous decay noise at late times.

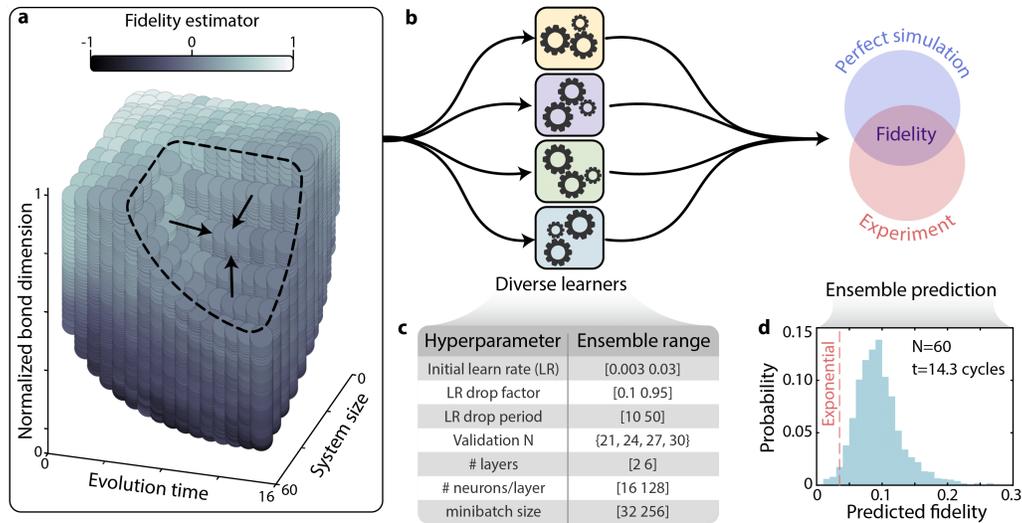


Figure 8.14: Monte Carlo fidelity inference. **a.** All the data in this study, mapping the three inputs (system size, N , evolution time, t , and log-normalized bond dimension, $\tilde{\chi} = \log(\chi)/\log(\chi_0)$) to an estimated fidelity. Here χ_0 is the time- and size-dependent saturation bond dimension (Fig. 8.17e). **b.** To improve generalization into the unknown regime, we use a Monte Carlo inference approach. We train an ensemble of 1500 fully-connected neural networks, each instantiated and trained with random choices of parameters within reasonably specified ranges (defined in **c**); the Monte Carlo predicted fidelity is then taken as the ensemble average for a given input. **d.** Monte Carlo prediction for the fidelity of the $N = 60$ system at the entanglement saturation time. Over $\sim 98\%$ of the individual Monte Carlo instance predictions exceed the expectation from the early time exponential decay (see Fig. 8.3), resulting in an average of $0.095(11)$, where the error bar represents the standard error on the mean, added in quadrature with the intrinsic sampling error of the underlying data.

8.5 MONTE CARLO INFERENCE

To extrapolate fidelities into the high-entanglement (i.e. late-time, large system-size) regime, we use a Monte Carlo inference approach. More specifically, we train an ensemble [397] of fully-connected regression neural networks, implemented in MATLAB, with three inputs (system size, evolution time, and bond dimension), and one output (estimated fidelity). The neural networks are instantiated with randomized hyperparameters (including the initial seed for the weights and biases). We emphasize that this approach is not trying to learn directly from experimental observations [405–408], and instead we are simply extrapolating a smoothly varying 3-to-1 function. Fundamentally, the goal of this approach is to utilize all aspects of the data simultaneously, rather than extrapolating along only a single axis at a time, for instance either via time-extrapolation or system-size scaling alone.

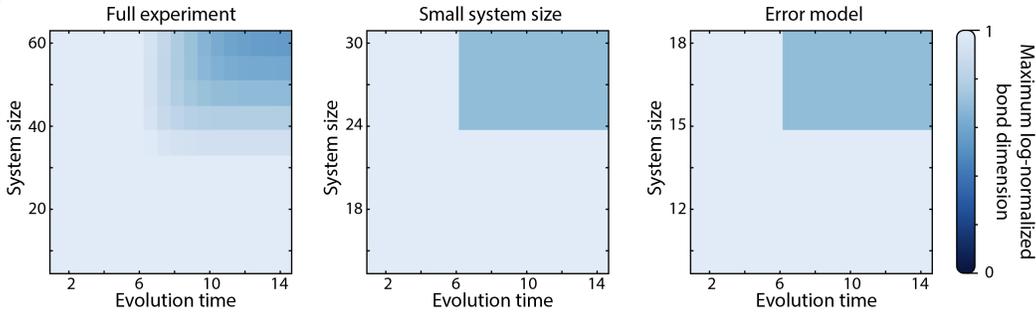


Figure 8.15: **Bond dimension training limits.** Maximum log-normalized bond dimension available for training as a function of system size and time, for $\chi = 3072$.

8.5.1 Method description

For training data, we use the fidelities from all combinations of the 14 system sizes, 18 evolution times, and 91 simulated bond dimensions (from $\chi = 2$ to 3072), resulting in a total dataset of ~ 23000 points (Fig. 8.14a). The network is then trained to learn the functional dependence of $F(N, \chi, t)$.

To improve generalization and avoid fine-tuning, we employ a Monte Carlo approach through ensemble learning [397], where the training is repeated 1500 times. Each iteration is instantiated with a randomized seed, and every three iterations we randomly select from a range of hyperparameters, the domain of which are specified in Fig. 8.14c. The fidelity for a given set of inputs is then taken to be the unweighted ensemble average, while the error bar is taken as the ensemble standard error of the mean added in quadrature with the sampling error at a given time.

To preprocess the data, we log-normalize the bond dimensions as $\tilde{\chi} = \log(\chi)/\log(\chi_0)$, where χ_0 is the size- and time-dependent saturation bond dimension required for the simulation to be exact (Fig. 8.17e). This makes the learned function have a more consistent domain for different system sizes and evolution times. We also test if we train using simply $\log(\chi)$. We stress that in doing so, the largest trained $\log(\chi) \approx 8$, but for instance for $N = 60$ we predict the fidelity at $\log(\chi) \approx 14.5$, with no training data in the interim regime. We find consistent mean predictions between the two methods, but the standard deviation of Monte Carlo instances is reduced by training on $\tilde{\chi}$.

We determine χ_0 by finding the minimum χ , at a given system size and evolution time, where the MPS fidelity is greater than 0.999. We extrapolate the behavior of χ_0 into the high-entanglement regime of large system sizes and late evolution times (Fig. 8.17e). We further normalize the input (subtracting off the mean and dividing by the standard deviation) to improve learnability.

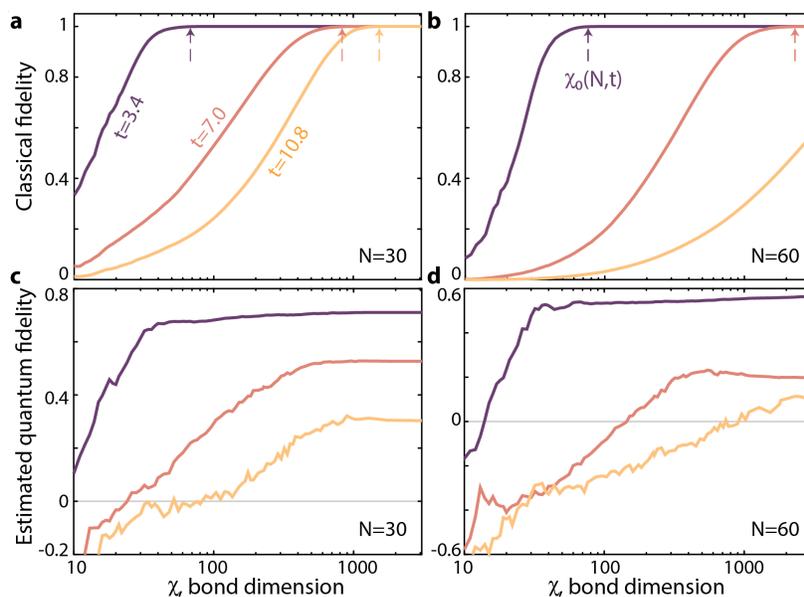


Figure 8.16: **Determining the log-normalized bond dimension.** **a, b.** As a function of bond dimension, the classical simulation fidelity rises before saturating to 1; we define χ_0 (the saturation bond dimension), as the minimum value for which the classical fidelity is greater than 0.999. **c, d.** By χ_0 , the experimentally benchmarked fidelities have also saturated.

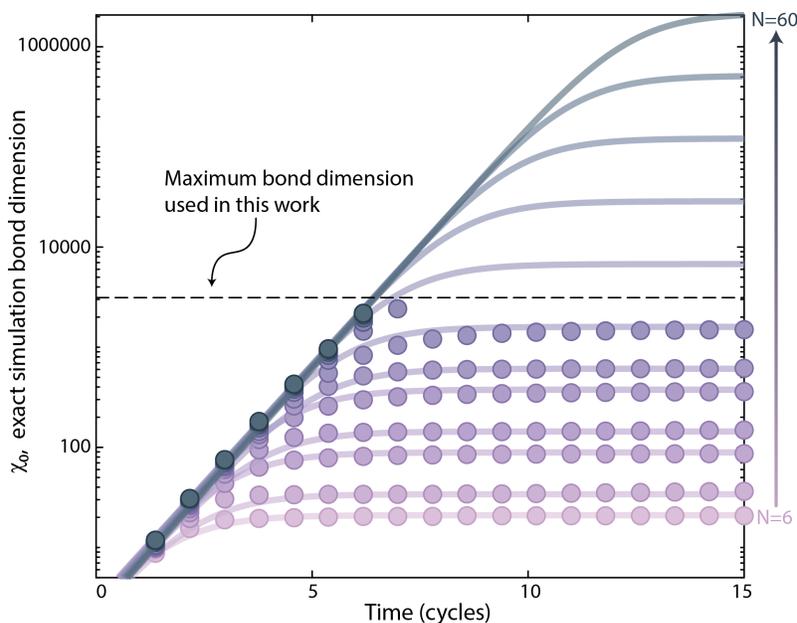


Figure 8.17: **Limits in log-normalized bond dimension.** We fit the dependence of χ_0 as a function of system size and time (solid lines) in order to extrapolate beyond the range of χ values which we directly simulate in this work.

To avoid overfitting, we randomly select a few N from 21 to 30, for which we excise from the training any data with $\tilde{\chi} > 2/3$ and $t > 6.6$ cycles. We then use the data for these system sizes and times with $\tilde{\chi} \sim 1$ as the validation dataset, and stop the training when the validation loss – defined by the prediction root mean square error (RMSE) – does not decrease for 50 epochs in a row (see for instance Fig. 8.19b).

We summarize a partial subset of the data showing the Monte Carlo inference as a function of N , t , and χ in Fig. 8.18.

8.5.2 Consistency checks

In order to test that this approach is accurate in estimating the fidelity, we perform several consistency checks, ranging from varying the number of learnable parameters, testing the result as a function of the amount of experimental training data, checking against small system sizes in experiment and error model where we have exact ground truth data, and more. Several of these tests are described below.

Small system sizes, experimental

For the largest system sizes we do not have ground truth data for F_d in the late-time regime, and so cannot directly verify our model’s efficacy. However, we can emulate the effect in this regime at smaller system sizes where we do have ground truth data in all regimes.

We employ the same training methodology as on the full dataset, but only consider $N \leq 30$ where we have the full ground truth. For $t > 6.6$ cycles and $N \geq 24$, we excise from the training data all bond dimensions for which $\tilde{\chi} > 2/3$. This value is chosen as it is approximately the maximum normalized bond dimension available at $N = 60$ (Fig. 8.15e). Concretely, this approach is essentially pretending as if $N \geq 24$ is beyond the exactly simulable regime. In reality, of course, we have the ground truth data to compare against, and can for instance use it to study the behavior of the loss functions (Fig. 8.19ab).

Averaging over all Monte Carlo instantiations, we find the training loss decreases consistently as a function of training epoch, but the validation loss (where the ground truth $N \geq 24$ results are used as the validation data set) reaches a minimum after ~ 50 epochs, after which it proceeds into an overfitting regime. As with the full dataset, we thus stop the training for each instantiation at the point where the validation loss is non-decreasing for ~ 50 epochs in a row, and then take the result from the best epoch. We note that both losses drop most dramatically after only a single epoch of training from initially randomized configurations, meaning the gradient descent quickly finds a nearly optimal regime, implying that the loss function is both smooth and non-flat.

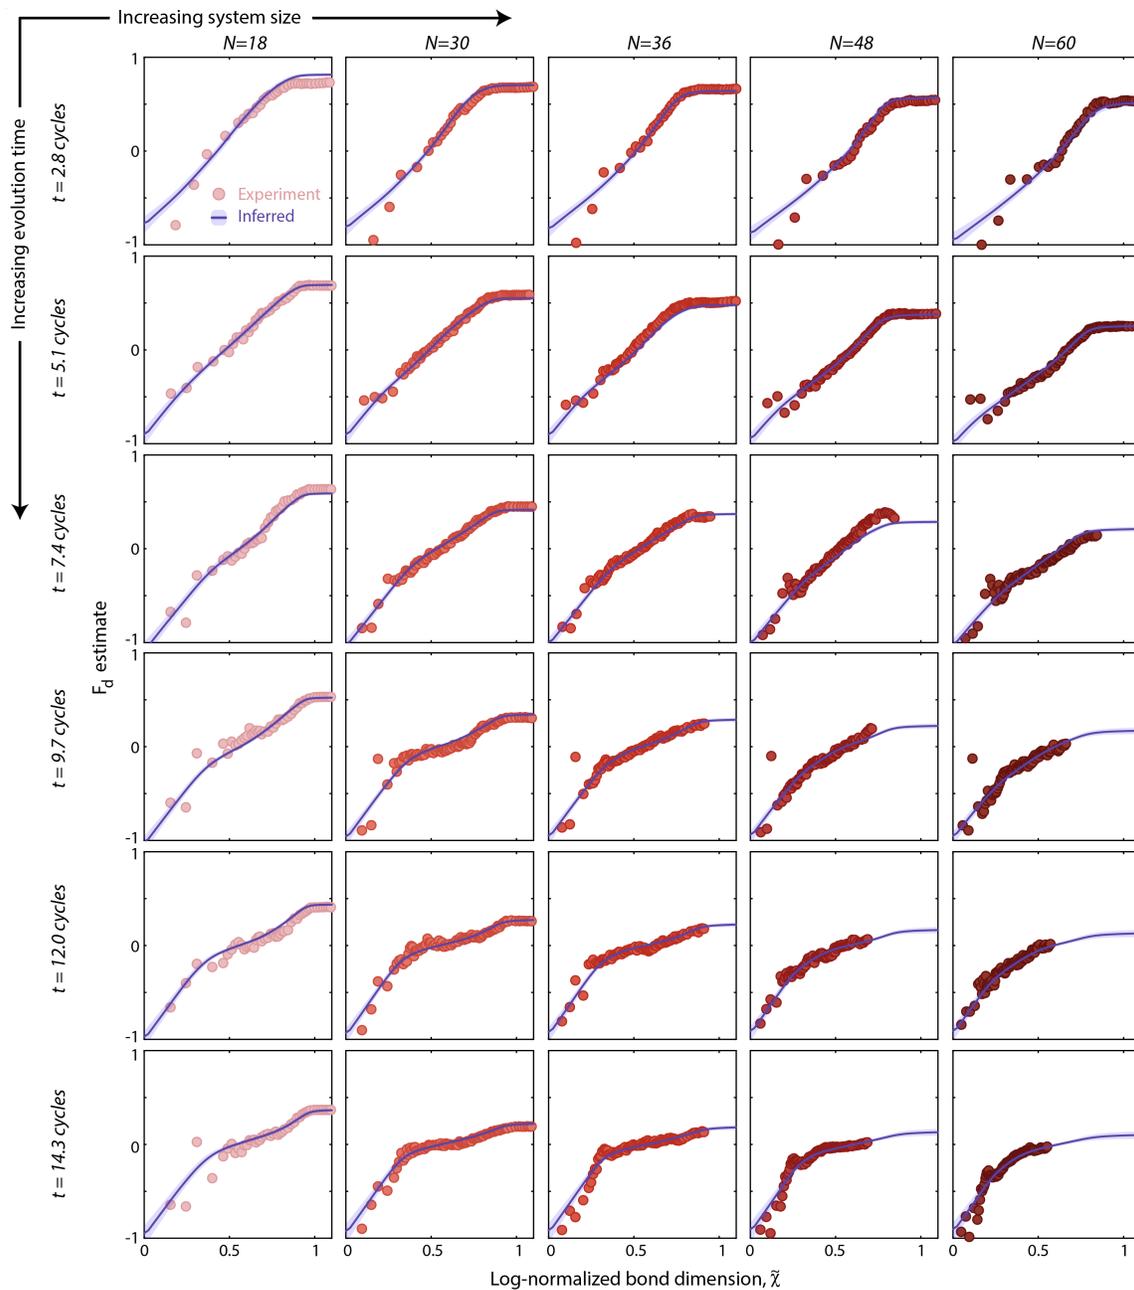


Figure 8.18: **Monte Carlo inference for all data.** **a.** Fidelity-estimator, F_d , for various system sizes and evolution times as a function of the log-normalized bond dimension, $\tilde{\chi} = \log(\chi)/\log(\chi_0)$. The fidelity is estimated at $\tilde{\chi} = 1$. A key feature to note is that as system size or evolution time increases (i.e. as χ_0 increases), the data points shift left: the classical simulation breakdown regime is when no data is available at $\tilde{\chi} = 1$.

We then compare directly between the Monte Carlo predicted F_d and the ground truth value for $N = 30$ across all times. We see good agreement even in the emulated breakdown regime at long times (Fig. 8.19c), with no indication that the learned fidelities are misestimating

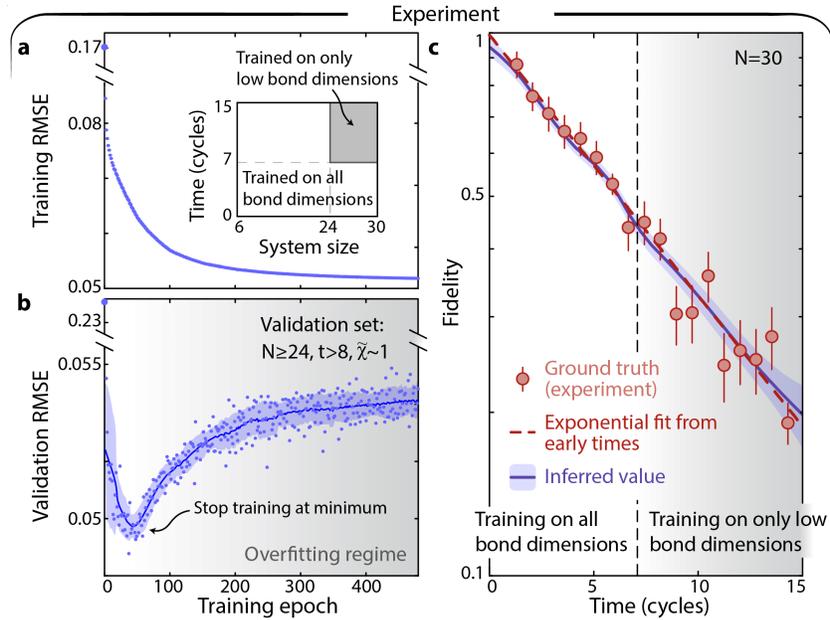


Figure 8.19: **Evaluating the Monte Carlo inference at small system sizes.** **a, b.** As a cross-check of our approach, we use the same training methodology employed on the full dataset, but truncate to only include $N \leq 30$ where we have the full ground truth. We bipartition the data (inset), and for the large system sizes and late evolution times we only directly train with $\tilde{\chi} < 2/3$. Averaging over 1500 instantiations, we find the training loss decreases consistently, but validation loss at $\tilde{\chi} \sim 1$ reaches a minimum after ~ 50 epochs, and as expected proceeds into an over-fitting regime after. **c.** Model predictions show the Monte Carlo inference accurately predicts fidelities at late times, despite not having access to the underlying ground truth data at large bond dimensions in this regime.

the true values, despite not having access to the underlying ground truth data at large bond dimensions.

Small system sizes, error model

As described, we can check our model's efficacy at small system sizes directly on experimental data, but the non-exponential fidelity decay behavior is not yet strongly visible. We reach this regime using Monte Carlo fidelity inference on error model simulations (Fig. 8.3). Here we further describe this analysis.

We perform error model simulations for $N = 8 \rightarrow 18$ up to the maximum experimental time, but artificially increase the shot-to-shot Rabi frequency variation by a factor of 4 in order to roughly emulate the predicted experimental fidelities at large N . We choose to do this, rather than simulate for much longer timescales where the behavior would become naturally evident (Fig. 8.11d), because we want to replicate as much as possible the fact that for the experimental dataset the entanglement is still growing for most times for the largest

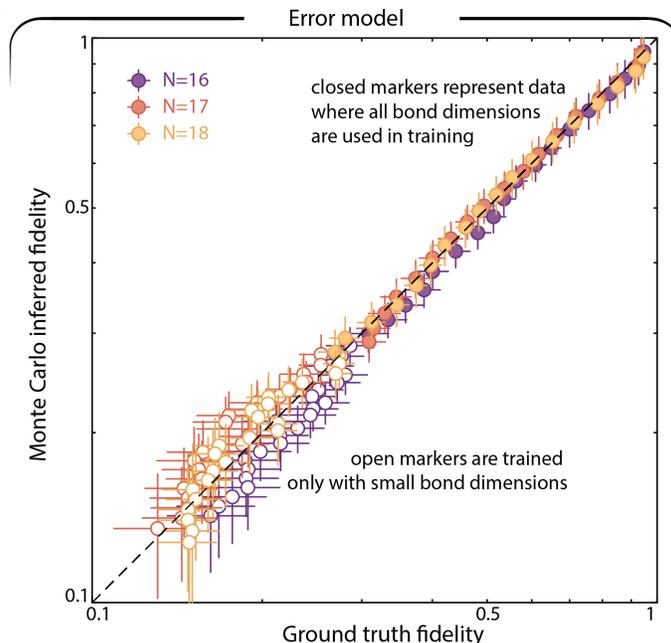


Figure 8.20: **Evaluating the Monte Carlo inference at small system sizes.** We perform a similar test using error model simulations of our experiment, as for instance shown in Fig. 8.3b. We train with $N = 8$ to 18, excising $\tilde{\chi} > 2/3$ for $N > 15$ and $t > 6.6$ cycles. The resultant Monte Carlo inference predictions are well correlated with the ground truth, across the whole range of fidelities. X-error bars represent estimated typicality fluctuations, y-error bars represent the Monte Carlo inference standard deviation.

system sizes. Performing error model simulations out to ~ 100 cycles would dominate the training by a regime where the entanglement had already saturated for these small system sizes.

We perform MPS simulations for all bond dimensions up to $\chi_0(N = 18, t \rightarrow \infty) = 89$, and benchmark the error model with each. Though we run the error model with a finer timestep of ~ 10 ns, we only select data at the experimentally measured times for training. We then emulate the experimental dataset by treating all $N < 16$ as “small N” where we have full ground truth data. We excise from the training any data for which $\tilde{\chi} > 2/3$, $N > 15$, and $t > 6.6$ cycles, to emulate the available data for the full experimental dataset. We use $N = 13 \rightarrow 15$ as the validation dataset to prevent overfitting. We then train in the same way as for the experiment, and compare Monte Carlo predictions versus ground truth output (Fig. 8.20). We find the two are consistent, both in terms of the values obtained, and the non-exponential fidelity decay, indicating the Monte Carlo inference is able to learn this behavior.

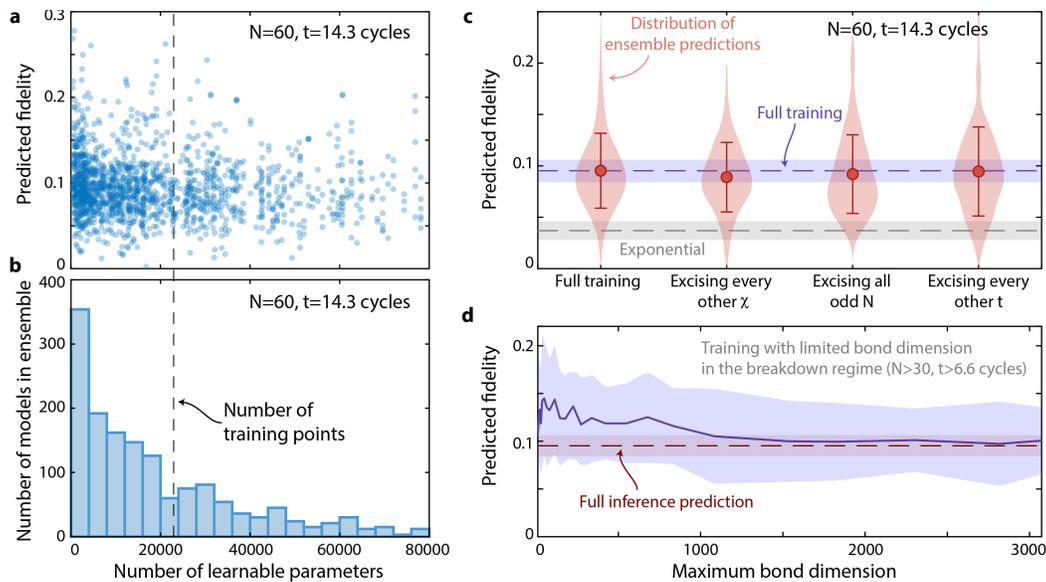


Figure 8.21: **Robustness of Monte Carlo inference.** **a, b.** For the majority ($\sim 70\%$) of hyperparameter combinations sampled during the Monte Carlo inference (Fig. 8.14c), the number of learnable parameters is smaller than the number of training points (~ 23000). Further, we see no correlation between number of learnable parameters (nor any individual specific choice of hyperparameter) with the predicted fidelity. **c.** The Monte Carlo mean prediction is insensitive to the choice of data: removing half of the data along any axis leads to the same predicted fidelity. Note that while the “Full training” is composed of 1500 models, the “excised” ensembles only feature 192 trained models. Markers show mean and standard deviation for each distribution. **d.** To check the convergence of our protocol, we repeat the Monte Carlo inference procedure with varying maximum bond dimension allowed in the breakdown regime of large system sizes and long evolution times. We find the fidelity (at $N = 60$ and the latest experimental time) is generally consistent over the entire range, with convergence achieved around a bond dimension of 1500.

Number of learnable parameters

For the full dataset, a total of ~ 23000 training data points are generated, while depending on the choice of hyperparameters, the number of learnable parameters can vary from $\sim 350 - 83000$. For $\sim 70\%$ of the Monte Carlo instances, the number of learnable parameters is lower than the number of training points (Fig. 8.21ab). Further we do not see any strong correlation between the predicted fidelity and the number of learnable parameters, or more generally between the fidelity and either network depth or width.

We also perform the same training procedure on several variants of reduced experimental datasets, where we excise from the training half of the bond dimensions, system sizes, or evolution times used. We find the resulting Monte Carlo mean is always consistent with the prediction from training on the full dataset (Fig. 8.21c).

Maximum bond dimension used

The maximum bond dimension used in the training is constrained by our computational resources, but we can still study if the Monte Carlo prediction appears to be converged as a function of the bond dimension. In Fig. 8.21d we systematically scale the training bond dimension used in the breakdown regime of large system sizes and long evolution times, and find that the results appear converged with a bond dimension of ~ 1500 , about a factor of 2 lower than the maximum bond dimension we consider in this work. However, we note that by supplementing the experimental data with numerical data obtained from our error model simulations, this required bond dimension is dramatically reduced (see the next section, Fig. 8.22).

8.6 SCALABILITY OF FIDELITY ESTIMATION

Importantly, our method of benchmarking large scale quantum systems from many approximate classical simulations is not yet at its limits. Indeed, we predict that it can be applied for much larger systems, especially making assumptions of the noise affecting the system.

8.6.1 Markovian noise

A major advantage of the fidelity estimator which we use, F_d , is its accuracy at early times [29, 30]. This is in contrast to the original linear cross-entropy which only becomes accurate for deep circuits [162]. As we established, using tensor network simulation methods such as MPS, we are able to compute exact classical simulation references for essentially arbitrarily large systems up to constant time. Assuming the noise affecting the system is purely Markovian (i.e. the fidelity decays exponentially [348]), then to estimate the fidelity in the late-time, high-entanglement regime it is sufficient to measure the fidelity with high-precision at early times, and then fit the resulting exponential decay rate.

This gives another axis for fidelity extrapolation which may be particularly useful for digital quantum devices implementing random unitary circuits, which typically use a form of system size scaling to estimate the fidelity in the high-entanglement regime [316]. Such an approach should work for fidelity estimation of systems with largely arbitrary size and dimension, but does require the exact functional form of the fidelity decay be assumed *a priori*, and thus would not apply for systems with more non-trivial noise sources like those affecting our present experiment (or potentially also some digital quantum computers [162]).

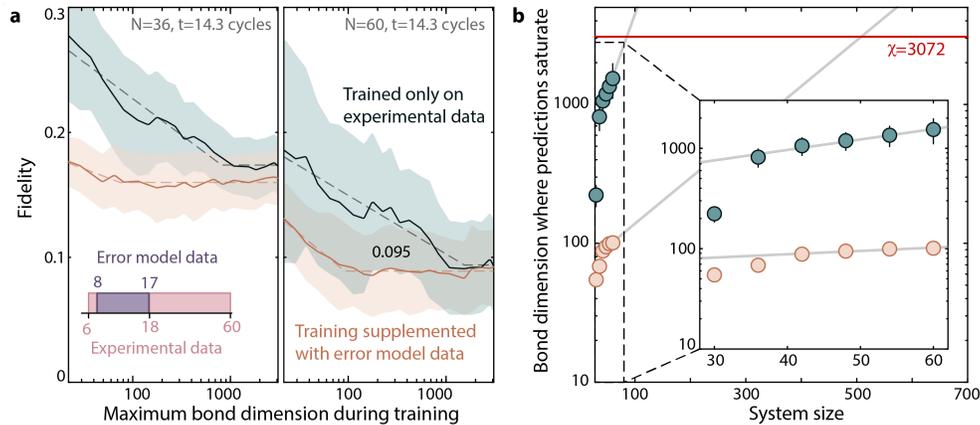


Figure 8.22: Scalability of Monte Carlo inference. **a.** We plot the predicted fidelity for $N = 36$ and $N = 60$ at the latest experimental time for two possible sets of training data: 1) considering only the experimental data (from $N = 6$ to $N = 60$ up to 15 cycles), and 2) supplementing with error model data out to late times (75 cycles). In order to improve scalability of our Monte Carlo inference protocol, we supplement training with error model data evaluated out to late evolution times for system sizes from $N = 8$ and $N = 17$. In both cases we vary the maximum bond dimension used during training, and then fit to find the bond dimension where predictions saturate, χ_{extrap} . Fits are proportional to $\log(\chi)$ and then constant after χ_{extrap} (dashed lines). **b.** χ_{extrap} as a function of system size for the predicted fidelity to saturate at the latest experimental time. The χ_{extrap} is improved by roughly an order of magnitude by supplementing with error model data. Conservatively extrapolating with $\chi_{\text{extrap}} \propto \exp(N)$ (grey lines) and setting a maximum $\chi = 3072$ shows maximum system sizes of ~ 500 and ~ 90 with and without the error model supplement.

8.6.2 Non-Markovian noise

In Fig. 8.21d, we investigated the minimum bond dimension of the classical (approximate) simulation, χ_{extrap} , required for the Monte Carlo inference procedure to reach a reliable, saturated prediction for the fidelity at $N = 60$ at the latest experimental time (14.5 cycles). There we found χ_{extrap} saturated for a bond dimension around $\chi \sim 1500$. However, an important question concerns the applicability at the Monte Carlo inference technique at larger system sizes, i.e. the scaling of χ_{extrap} with N . If the required bond dimension scales poorly with system size then it becomes difficult to trust the Monte Carlo inference results at larger system sizes, assuming there is a cap in the maximum bond dimension usable by the classical simulation.

We consider two versions of the training procedure (Fig. 8.22). In the first, we perform Monte Carlo inference using just the experimental data, as done previously. In the second approach we replace the training data from $N = 8$ to 17 with data obtained from error model simulations up to 75 evolution cycles (sampled every 1 cycle). F_d values are calculated for

error model data for χ ranging from 1 to 90. We replace with error model results, rather than simply augmenting, so that the training data does not have multiple outputs defined for the same set of inputs.

We have extensively checked the validity of our error model both in terms of fidelity and KL divergence (also see Refs. [26, 29] and Fig. 8.6), and using error model allows us to generate training data for F_d as a function of bond dimension in the low- N , long-time regime where the non-exponential fidelity decay becomes more clearly visible. We emphasize that access to an error model is not a general requirement, but is instead simply a tool we are employing in this case as the original experimental dataset was planned and taken without accounting for the non-exponential fidelity decay, as it was a novel observation of our present analysis. Using the error model alone (with no experimental data) for inference is not expected to give good fidelity extrapolation at large system sizes given the large difference between the largest simulable error model size ($N \approx 20$) and the largest experimental size.

In both cases (purely experimental data, or supplemented with error model data), we plot χ_{extrap} as a function of system size. Here χ_{extrap} is found from a phenomenological fit of the predicted fidelity, F_p , as a function of the maximum bond dimension used in training, χ ,

$$F_p = \begin{cases} F + A \log(\chi) & \chi < \chi_{\text{extrap}} \\ F & \chi \geq \chi_{\text{extrap}}, \end{cases} \quad (8.21)$$

where F and A are free fit parameters (Fig. 8.22c). In general, we see that supplementing with error model data produces consistent fidelity predictions as our original, all-experimental inferences, but saturates far earlier. This builds additional confidence in both the error model and the Monte Carlo inference procedure.

We then study how χ_{extrap} scales as a function of N , and set a cap of $\chi \sim 3072$, which represents a significant, but still achievable bond dimension. We choose a conservative phenomenological fit model of

$$\chi_{\text{extrap}} = A \exp(BN), \quad (8.22)$$

where A and B are fit constants.

We find the all-experimental inference exceeds this cap for $N \sim 90$, but the supplemented dataset goes as high as $N \sim 500$. We emphasize that this value represents a lower bound that could be improved either through use of larger classical resources, or possibly through an improved classical algorithm. This indicates that as long as care is taken to acquire experimental data in a diverse regime of system sizes and evolution times, we conserva-

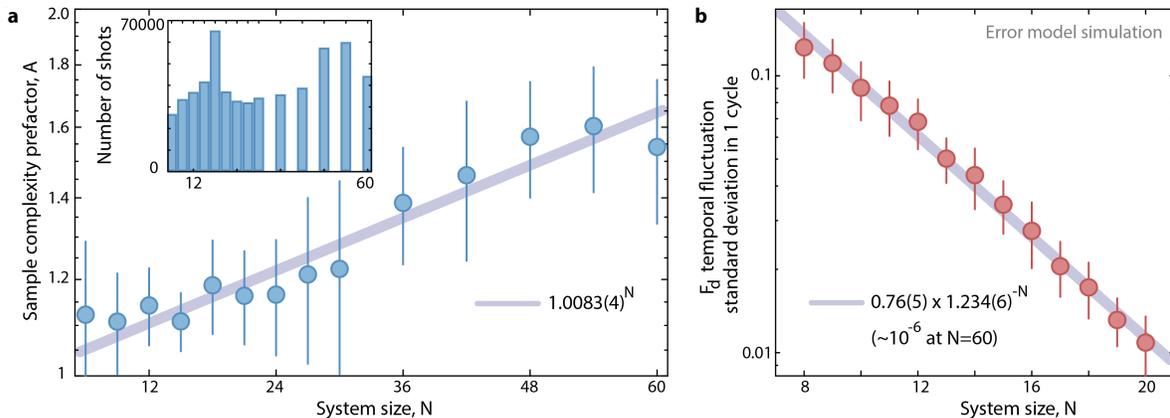


Figure 8.23: Uncertainty sources of the benchmarking protocol. There are two fundamental sources of noise in calculating F_d : statistical error, and systematic error arising from so-called typicality fluctuations [29]. **a.** We first evaluate the statistical error behavior. At a fixed fidelity, the statistical error bars of our protocol are given [30] by A/\sqrt{M} , where M is the number of samples and A is the sample complexity. At approximately a fixed fidelity of $F_d = 0.5$ (see text), we find A grows weakly with system size, indicating favorable scalability of our protocol from a statistical perspective. For this work, $M \sim 40000$ for all system sizes, summing over all times (inset). **b.** Typicality error manifests as temporal fluctuations of the estimator F_d , with respect to the true fidelity F , due to benchmarking in a finite-sized Hilbert space. The standard deviation of these fluctuations scales [29] as $1/\sqrt{D}$, where D is the Hilbert space dimension. From error model simulations, we calculate the standard deviation of temporal fluctuations of F_d in a 1-cycle period at late times, which we find decreases as 1.23^{-N} (or $\sim 1.2/\sqrt{D}$), implying a negligible typicality error of $\sim 10^{-6}$ at $N = 60$.

tively believe the Monte Carlo inference procedure we show here should provide scalable predictions for an order of magnitude more atoms than we currently study.

8.6.3 Sample complexity

To estimate error bars on the raw data, σ_F , we propagate from separately calculating the sampling error of the numerator and denominator of F_d as

$$\sigma = \sqrt{(\mathbb{E}(\tilde{p}(z)^2) - \mathbb{E}(\tilde{p}(z))^2) / M}, \quad (8.23)$$

where \tilde{p} is the rescaled bitstring probability distribution obtained from MPS simulation, M is the number of samples, and the expectation value is taken over the set of experimentally (numerically) sampled bitstrings for the numerator (denominator). Enough samples are taken numerically that the dominant overall sampling error is from the experiment. See Section B for more details.

We expect the statistical error of our benchmarking protocol [29, 30] to scale as A/\sqrt{M} , where A is the *sample complexity* which scales exponentially in system size with a base which depends on the fidelity [30]. We measure A at an approximately fixed fidelity of $F_d = 0.5$ by performing an average of all the experimentally measured error bars from Eq. (8.23), weighted by their inverse distance from $F_d = 0.5$, such that

$$A(N) = \frac{\sum_t \sigma_F(N, t) \sqrt{M(N, t)} w(N, t)}{\sum_t w(N, t)}, \quad (8.24)$$

with $w(N, t) = 1/|F_d(N, t) - 0.5|$. With this approach, we find $A = 1.0083(4)^N$ (Fig. 8.23a).

Quantum comparisons

Having benchmarked the fidelity of our Rydberg quantum simulator, we now turn to investigate the actual half-chain bipartite entanglement content of the experiment. In the past, several studies have investigated entanglement properties of (nearly) pure states by estimating the second Rényi entropy in (sub-)systems up to 10 particles [306, 409–411]. However, the actual output of an experiment can be a highly mixed state, with markedly different entanglement content than the target pure state. For this reason, it is desirable to directly quantify *mixed state entanglement* measures. Unfortunately, extensions of most pure state entanglement measures to the case of mixed states are defined variationally, and as such are in calculable for even moderately sized systems [412]. For a general review and tutorial for concepts related to entanglement entropy, see Ref. [413]

Here we devise and demonstrate a new proxy for the mixed state entanglement of a system, based on the negativity [72]. Our proxy takes as input only the (assumed or extrapolated) pure state entanglement, and the many-body fidelity. As such, we are able to calculate it whenever fidelity-estimation is possible. This estimator has several important applications, but in particular extends the ideas of Section 6.3.3 as a universal proxy comparable across all quantum platforms.

Our discussion is broken into four parts.

First, we introduce our entanglement proxy, demonstrate its effectiveness for numerical simulations of both digital and analog quantum systems, and then show a comparison amongst many state-of-the-art quantum platforms.

Next, we prove certain bounds about the effectiveness of our proxy. In particular, we prove its validity for the case when the target pure state, $|\psi\rangle$ is an eigenstate of the mixed state, $\hat{\rho}$,

as for instance is the case if $|\psi\rangle$ is the highest fidelity pure state to $\hat{\rho}$. We further improve our result for the case where $|\psi\rangle$ is a Haar-random state. Finally, we prove more general unconditional upper and lower bounds to argue that our entanglement-proxy becomes exponentially tighter as the system size increases.

Then, we discuss possible violations of the lower bound given by \mathcal{E}_P , if we allow $|\psi\rangle$ to no longer be an eigenstate of $\hat{\rho}$. We consider two physically realistic possibilities: 1) globally correlated coherent errors, and 2) incoherent local errors. For each case we provide numeric or analytic evidence showing that at worst our entanglement-proxy lower bound is potentially violated by an $\mathcal{O}(1)$ amount.

Finally, we show the entanglement-proxy is more broadly applicable than solely for studying quench dynamics, and for instance can be used to estimate entanglement build-up during quasiadiabatic state preparation. Enabling this application is the work of *target state benchmarking*, as discussed in Section 6.3.2.

9.1 EXPERIMENTAL MIXED STATE ENTANGLEMENT

While pure-state entanglement measures have often been estimated experimentally [306, 409–411], extensions of most pure-state entanglement measures to the case of *mixed states* are defined variationally, and as such are incalculable for even moderately sized systems [412]. An alternative, computable measure of mixed state entanglement is the log negativity [72], \mathcal{E}_N , which is an upper bound to the distillable entanglement of the system [412]. However, measuring the value of the negativity naïvely requires tomography of the full system density matrix, which is infeasible even for intermediate scale quantum systems [313, 314]. In the past, experiments have been limited to demonstrating necessary conditions for a non-vanishing negativity, which can only reveal the binary presence of mixed state entanglement [414, 415].

Explicitly, consider a system bipartitioned into two subsystems, A and B , and described by density matrix, $\hat{\rho}_{AB}$. The log negativity [72] is given by

$$\mathcal{E}_N = \log_2 \|\hat{\rho}_{AB}^{T_A}\|_1,$$

where $\hat{\rho}_{AB}^{T_A}$ is the partially transposed density matrix, and $\|\cdot\|_1$ is the trace norm. More plainly, the log negativity measures the log absolute sum of all eigenvalues of $\hat{\rho}_{AB}^{T_A}$. If $\hat{\rho}_{AB}$ is separable, the log negativity is zero, but for entangled states it generically serves as an upper bound to the distillable entanglement of the system [412], and is an entanglement monotone [412]. For pure states the log negativity is equivalent to the Rényi-1/2 entropy.

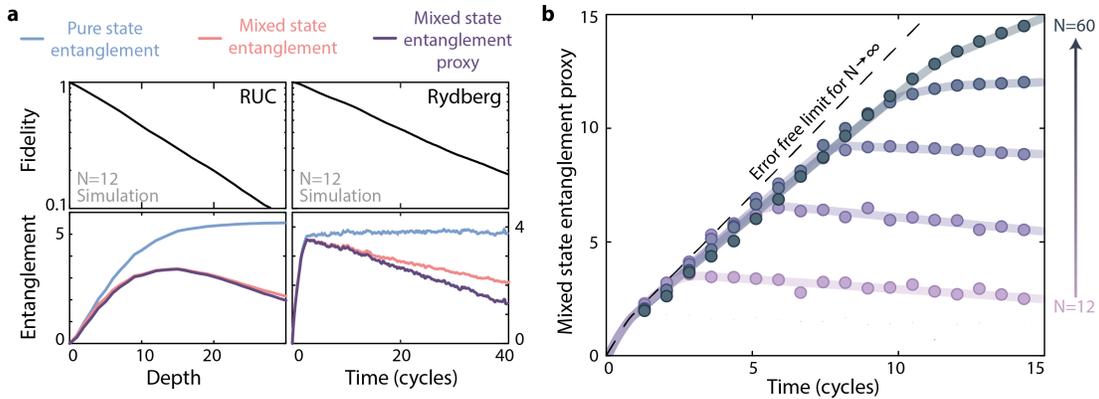


Figure 9.1: **Experimental mixed state entanglement.** **a.** We develop an experimentally measurable proxy that lower-bounds the log negativity, which is a measure of mixed state entanglement. Here we demonstrate this proxy with error model simulations of random unitary circuit and Rydberg evolution. **b.** The experimental mixed state entanglement-proxy; solid lines are guides to the eye.

Here we derive and demonstrate an entanglement-proxy, \mathcal{E}_P , which can lower-bound the extensive mixed state entanglement (quantified by log negativity). For a mixed state, $\hat{\rho}$, with fidelity, F , to a target pure state, $|\psi\rangle$, with known entanglement, $\mathcal{E}_N(|\psi\rangle)$, our mixed state entanglement-proxy is

$$\mathcal{E}_P(\hat{\rho}) \equiv \mathcal{E}_N(|\psi\rangle) + \log_2(F). \quad (9.1)$$

Intuitively, \mathcal{E}_P is a proxy evaluating the competition between the growth of the error-free entanglement, $\mathcal{E}_N(|\psi\rangle)$, versus the error-sensitive fidelity, as $F < 1$ reduces the mixed state entanglement. When $\hat{\rho}$ is an isotropic state (an admixture of a maximally entangled state and a maximally mixed state), it has been shown [72, 416] that $\mathcal{E}_N(\hat{\rho}) = \max(\mathcal{E}_P(\hat{\rho}), 0)$ at large system sizes. Further, we show the same holds for a Haar-random state admixed with a maximally mixed state – the expected output [348] of deep noisy random unitary circuits (RUCs) – as long as the fidelity is large compared to the inverse of the half-chain Hilbert space dimension.

More generally, we prove \mathcal{E}_P is a lower bound for \mathcal{E}_N for *any* mixed state assuming $|\psi\rangle$ is the highest fidelity state to $\hat{\rho}$, and becomes tighter as the system size increases (Fig. 9.4). Importantly, violations of this assumption can only lead to small violations of our bound in the worst case for physically realistic conditions with local or quasi-static errors, as we show with both analytic (Theorems 3 and 4) and numeric (Figs. 9.5 and 9.8) support.

We demonstrate the efficacy of \mathcal{E}_P on both noisy RUC evolution and error model simulation of our Rydberg dynamics (Fig. 9.1a). In both cases, the target pure state log negativity

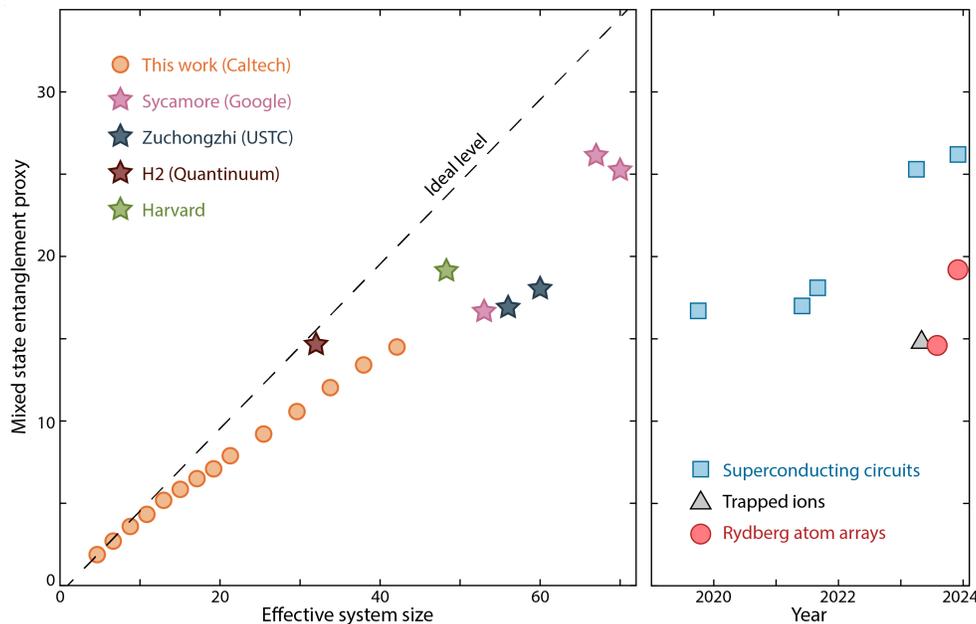


Figure 9.2: **Mixed state entanglement comparisons.** The maximum entanglement-proxy for our experiment can be compared against that of literature examples performing global fidelity estimation with digital quantum processors: Sycamore [162, 316], Zuchongzhi [381, 382], and H2 [395] (text indicates release year). For literature examples, the x-axis is the number of qubits, while for our experiment the effective system size is defined as the number of qubits with the same Hilbert space dimension as our experiment under the Rydberg blockade constraint, and is for instance ~ 42 at $N = 60$.

Experiment	Type	Date	N	Fidelity ($\times 10^{-2}$)	Mixed state entanglement	Ref.
Sycamore	Digital	Oct 2019	53	0.22	16.7	[162]
Zuchongzhi 2.0	Digital	Jun 2021	56	0.066	17.0	[381]
Zuchongzhi 2.1	Digital	Sep 2021	60	0.037	18.1	[382]
Sycamore	Digital	Apr 2023	70	0.17	25.3	[316]
Sycamore	Digital	Dec 2023	67	0.9*	26.2	[316]
Quantinuum H2	Digital	May 2023	32	57	14.7	[395]
Harvard	Digital	Dec 2023	48	9.7	19.2	[40]
This work	Analog	Aug 2023	42 [†]	9.5	14.6	-

Table 9.1: **Mixed state entanglement values for various experiments.** For all literature experiments, we assume the target pure state is as entangled as a typical Haar random state, while for this work we extrapolate the entanglement directly (Fig. 9.3). [†] indicates an approximate effective system size due to the Rydberg blockade constraint, see text.

increases and saturates, while the exactly calculated mixed state log negativity reaches a maximum before decaying at late times, behavior which the entanglement-proxy \mathcal{E}_P replicates as a lower bound.

We then plot the experimental entanglement-proxy (Fig. 9.1b), where $\mathcal{E}_N(|\psi\rangle)$ is extrapolated from small system sizes (Fig. 9.3), and F is found from Monte Carlo inference. We observe the entanglement proxy peaks before falling at late times; this peak value increases (Fig. 9.2c) as a function of effective system size defined as the number of qubits with the same Hilbert space dimension as our experiment under the Rydberg blockade constraint (~ 42 for $N=60$).

Importantly, with Eq. (9.1) we can directly compare the results of our present study against RUC evolution in state-of-the-art digital quantum devices [162, 381, 382, 395] (Fig. 9.2c). We find we are within ~ 2 ebits of early tests of quantum advantage [162] (an ebit is the entanglement of a two-qubit Bell state). For literature examples, we assume targeted states are Haar-random [417, 418], while for our experiment we conservatively use the extrapolated log negativity, which is ~ 2 ebits below the expectation for Haar-random states at the largest system sizes (Fig. 9.3).

The mixed-state entanglement-proxy \mathcal{E}_P can serve as a useful quality-factor of the ability for different experiments to produce highly entangled states, including preparation methods besides quench evolution such as quasi-adiabatic ground state preparation (Figs. 9.10 and 9.9), and could be a more widely applicable alternative to other measures, such as quantum volume [277], for directing efforts to improve NISQ-era quantum systems.

9.2 NEGATIVITY BOUNDS AND ESTIMATION

Here we discuss our mixed state entanglement proxy in more detail, and specifically provide rigorous bounds guaranteeing its efficacy in reasonable regimes.

9.2.1 Extrapolating pure state entanglement growth

Calculating the mixed state entanglement proxy requires we have an estimate of $\mathcal{E}_N(|\psi\rangle)$, the negativity of the target pure state. For RUC evolution of sufficient depth, we can assume this is equal to the value predicted for Haar random states [324], but cannot necessarily make the same assumptions for our experiment. Thus we first must estimate the pure state entanglement, even for system sizes for which we can no longer exactly simulate the dynamics. To do this, we extrapolate the scaling behavior from small system sizes.

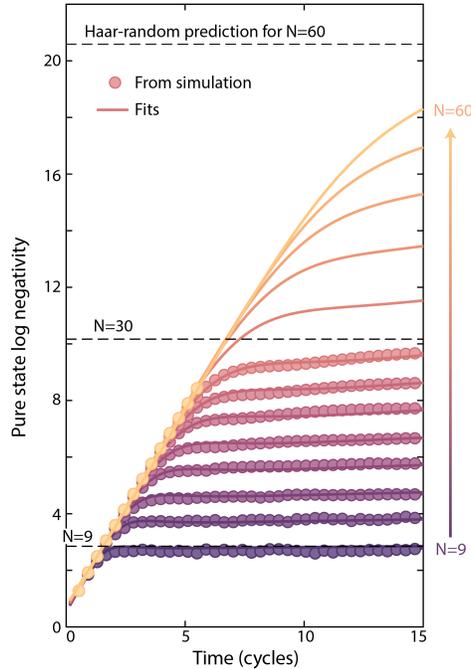


Figure 9.3: **Pure state entanglement growth dynamics for the Rydberg Hamiltonian.** We directly calculate the entanglement (here log negativity) for system sizes up to $N = 60$, and fit the resulting data (see text) to extrapolate the entanglement dynamics for large N at late times. We also compare against the prediction for Haar random states at various system sizes (black dashed lines), and find that as system size increases we gradually undershoot this expectation more, which we attribute potentially to imperfect parameter selection making the dynamics effectively less than infinite temperature at large N .

Specifically, we fit entanglement growth, S , with the following functional form as a function of time and system size:

$$S = S_0 + m_1 t + \frac{(m_2 - m_1)(t - t_{\text{ent}} - \sigma)}{1 + \exp(-\frac{t - t_{\text{ent}}}{\sigma})}. \quad (9.2)$$

To respect the physical constraints of the entanglement growth, we fix m_1 – the early time linear entanglement growth – to be system size independent. The Rydberg Hamiltonian induces fast entanglement growth over the first ~ 0.5 cycles of evolution which makes fitting the linear growth erroneous without also including a system-size independent offset, S_0 . We then expect t_{ent} , which we identify with the approximate entanglement saturation time, to scale linearly with N , which we empirically find is true also for σ (a smoothness parameter). We further empirically find that m_2 (the shallow entanglement growth slope after t_{ent}) scales approximately quadratically with N .

We show the results of these fits in Fig. 9.3, where markers are measured directly from

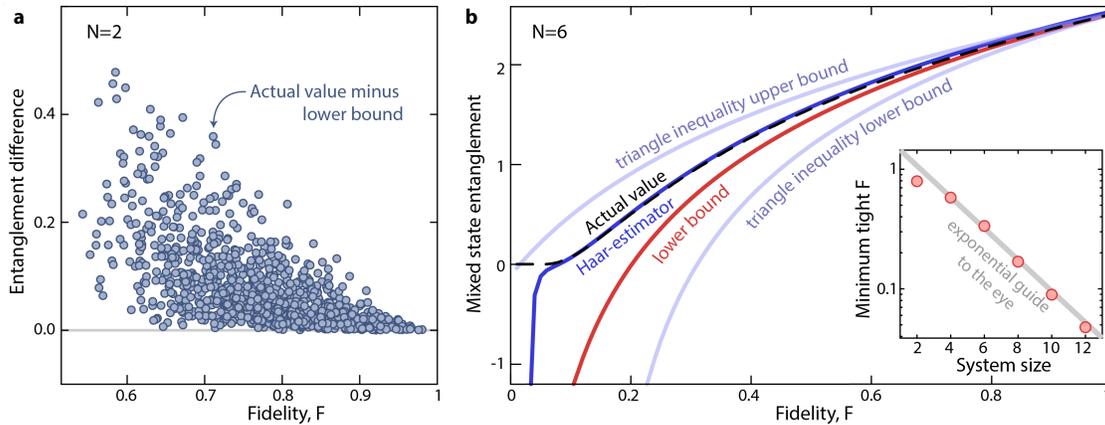


Figure 9.4: Estimation of log negativity from fidelity. **a.** We demonstrate the validity of the lower bound Eq. (9.1), by plotting the difference $\mathcal{E}_N(\hat{\rho}) - \mathcal{E}_P$, with $|\psi\rangle$ the pure state with highest fidelity to $\hat{\rho}$. Here, we generate 1000 random two-qubit mixed states as uniform incoherent mixtures of two Haar random states. The difference is always above zero, indicating the validity of the lower bound. **b.** Dependence between log negativity $\mathcal{E}_N(\hat{\rho})$ and fidelity F for globally depolarized Haar random states, Eq. (9.11). For small values of F , the negativity $\mathcal{E}_N(\hat{\rho})$ (black dashed) deviates from the lower bound \mathcal{E}_P (red), illustrated here for the half-chain negativity of six-qubit depolarized Haar random states. Remarkably, for such states, $\mathcal{E}_N(\hat{\rho})$ is uniquely determined by the fidelity F and Hilbert space dimension D . We find an analytic series expression for this dependence, Eq. (9.13), with first-order correction illustrated in blue. We also illustrate general lower and upper bounds for the log negativity, Eq. (9.18), in light blue. Inset: The lower bound (red) improves exponentially as a function of system size, quantified here as the minimum fidelity for which the lower bound is 90% of the actual value.

simulation, and lines are fits. Though here we show log negativity, the von Neumann entropy scales and acts in the same way.

It is important to note that as system size increases, we gradually reach relatively less entangled states at t_{ent} as compared to the expectation for a Haar random state (dashed lines in Fig. 9.3). Given the Rydberg blockade constraint, we define the Haar random prediction as $S_{\text{Haar}} \approx \log_2(D_0) - 0.47$, where D_0 is the blockaded Hilbert space dimension of the half-cut chain, and the numerical offset of -0.47 is a type of Page correction [324], which is actually $[417, 418] \approx \log_2(\frac{64}{9\pi^2})$, as explained further below.

9.2.2 Negativity lower bounds and estimators

The validity of our entanglement-proxy has been shown in the past [416] for the case of *isotropic states*; here we extend this result and validate our entanglement-proxy when the target pure state, $|\psi\rangle$, is an eigenstate of the mixed state, $\hat{\rho}$ (Theorem 2). We then improve

our result for the case of depolarized Haar-random pure states (Lemma 1), and then show more general lower bounds which allow us to argue our entanglement-proxy should tighten exponentially with system size (Lemma 2). Numerical support of these analytical claims is shown in Fig. 9.4.

We begin with the case of $|\psi\rangle$ being an eigenstate of $\hat{\rho}$

Theorem 2. *For any mixed state $\hat{\rho}$ and pure state $|\psi\rangle$ which is an eigenstate of $\hat{\rho}$, the logarithmic negativities $\mathcal{E}_N(\hat{\rho})$ and $\mathcal{E}_N(|\psi\rangle)$ are related by*

$$\mathcal{E}_N(\hat{\rho}) \geq \mathcal{E}_P \equiv \mathcal{E}_N(|\psi\rangle) + \log_2(F), \quad (9.3)$$

where $F \equiv \langle \psi | \hat{\rho} | \psi \rangle$ is their fidelity.

Proof. The proof of this Theorem follows in three steps. It is equivalent to derive the bound

$$\|\hat{\rho}^{TA}\|_1 \geq F \|\psi\rangle\langle\psi|^{TA}\|_1. \quad (9.4)$$

First, we express the state

$$\hat{\rho} = F|\psi\rangle\langle\psi| + (1 - F)\hat{\rho}_\perp \quad (9.5)$$

for some positive semi-definite state $\hat{\rho}_\perp$, which is guaranteed because $|\psi\rangle$ is an eigenstate of $\hat{\rho}$.

Next, we use the variational definition of the trace-norm as a maximum over projectors \hat{P} , i.e. operators with eigenvalues either 0 or 1.

$$\|\hat{\rho}^{TA}\|_1 = 2 \max_P \text{tr}(\hat{P}\hat{\rho}^{TA}) - 1. \quad (9.6)$$

This maximum is attained when \hat{P} is the projector onto the positive-eigenvalue subspace of $\hat{\rho}^{TA}$. This variational definition gives the lower bound:

$$\|\hat{\rho}^{TA}\|_1 \geq F \|\psi\rangle\langle\psi|^{TA}\|_1 + (1 - F)(2\text{tr}(\hat{P}_\psi\hat{\rho}_\perp^{TA}) - 1), \quad (9.7)$$

where \hat{P}_ψ is the projector onto the positive-eigenvalue subspace of $|\psi\rangle\langle\psi|^{TA}$.

Finally, we utilize an explicit construction of \hat{P}_ψ to show that $\text{tr}(\hat{P}_\psi\hat{\rho}_\perp^{TA}) \geq 1/2$, giving the desired bound Eq. (9.4). This step is the most involved, which we show below.

The eigenvalues λ and eigenvectors $|v\rangle$ of $|\psi\rangle\langle\psi|^{TA}$ can be expressed in terms of the Schmidt values s_j and Schmidt basis $|a_j\rangle, |b_j\rangle$ of $|\psi\rangle$, with respect to the same bipartition [416]. With $|\psi\rangle = \sum_j s_j |a_j\rangle |b_j\rangle$, the eigenvalues λ are s_j^2 or $\pm s_i s_j$ (for $i \neq j$), and their corresponding eigenvectors are respectively $|v_{jj}\rangle \equiv |a_j\rangle |b_j\rangle$ and $|v_{ij}^\pm\rangle \equiv \frac{1}{2} (|a_i\rangle |b_j\rangle \pm |a_j\rangle |b_i\rangle)$. Then \hat{P}_ψ has explicit expression:

$$\hat{P}_\psi = \sum_j |v_{jj}\rangle\langle v_{jj}| + \sum_{i<j} |v_{ij}^\pm\rangle\langle v_{ij}^\pm|. \quad (9.8)$$

Next, we use the relation $\text{tr}(\hat{M}^{TA} \hat{N}^{TA}) = \text{tr}(\hat{M} \hat{N})$, valid for all operators \hat{M} and \hat{N} defined on AB . Thus $\text{tr}(\hat{P}_\psi \hat{\rho}_\perp^{TA}) = \text{tr}(\hat{P}_\psi \hat{\rho}_\perp)$. \hat{P}_ψ^{TA} can be expressed as

$$\hat{P}_\psi^{TA} = \frac{1}{2} (I + |\Phi\rangle\langle\Phi|), \quad (9.9)$$

where $|\Phi\rangle \equiv \sum_j |a_j\rangle |b_j\rangle$ is the *unnormalized* maximally entangled state in the Schmidt bases $\{|a_j\rangle\}$ and $\{|b_j\rangle\}$. Since $\hat{\rho}_\perp$ is a positive semidefinite state, we have

$$\text{tr}(\hat{P}_\psi^{TA} \hat{\rho}_\perp) = \frac{1}{2} \text{tr}(\hat{\rho}_\perp) + \frac{1}{2} \langle\Phi| \hat{\rho}_\perp |\Phi\rangle \geq \frac{1}{2}. \quad (9.10)$$

This proves the desired bound Eq. (9.4).

This Theorem validates the entanglement proxy in the simplest case. We next move to improve it, first specifically in the case where $|\psi\rangle$ is taken to be a Haar random state, $|\psi_{\text{Haar}}\rangle$.

Lemma 1. *For depolarized Haar random states, i.e. states of the form*

$$\hat{\rho} = F' |\psi_{\text{Haar}}\rangle\langle\psi_{\text{Haar}}| + (1 - F') \frac{I}{D}, \quad (9.11)$$

where $F' = F - (1 - F)/(D - 1)$ (chosen so the fidelity $\langle\psi_{\text{Haar}}|\hat{\rho}|\psi_{\text{Haar}}\rangle$ equals F), the mixed state entanglement has the form

$$\mathcal{E}_N(\hat{\rho}) \approx \mathcal{E}_N(|\psi\rangle) + \log_2(\mathfrak{f}(F)), \quad (9.12)$$

where the more general function \mathfrak{f} can be expressed to low order as

$$\mathfrak{f}(F) = F' + \frac{9}{64} (F' D_A)^{-1} (2 \ln(F' D_A) + 8 \ln 2 - 1) + \mathcal{O}((F' D_A)^{-3}), \quad (9.13)$$

where D_A is the Hilbert space dimension of subsystem A , here assumed to be for a half-chain bipartitions.

Proof. Here, we restrict our discussion to equal bipartitions, although our results can be generalized to any bipartition.

First, with Haar random pure states, with high probability we have [417, 418]

$$\mathcal{E}_N(|\psi\rangle) = \frac{N}{2} + \log_2 \frac{64}{9\pi^2} \approx \frac{N}{2} - 0.472. \quad (9.14)$$

This expression follows from random matrix theory (RMT), which states that the Schmidt values s_j of a Haar random state follow the *quarter-circle law*, while s_j^2 follow the *Marchenko-Pastur distribution*.

This enables us to compute the distribution of eigenvalues λ of $|\psi\rangle\langle\psi|^{T_A}$. The eigenvalues of the form $\lambda=s_j^2$ contribute a constant amount to $\|\hat{\rho}^{T_A}\|_1$. The dominant (exponentially larger) contribution to the log negativity comes from the eigenvalues of the form $\lambda= \pm s_i s_j$ which follow the distribution [419]

$$P(\tilde{\lambda}) = \frac{2}{\pi^2} \left[\left(1 + \frac{\tilde{\lambda}^2}{16}\right) K \left(1 - \frac{\tilde{\lambda}^2}{16}\right) - 2E \left(1 - \frac{\tilde{\lambda}^2}{16}\right) \right], \quad (9.15)$$

where $\tilde{\lambda} \equiv D_A \lambda$, $P(\tilde{\lambda})$ is supported in $\tilde{\lambda} \in [-4, 4]$, and K and E are complete elliptic integrals.

The negativity of the depolarized state $\hat{\rho}$ can be calculated using this distribution: the eigenvalues λ' of $\hat{\rho}^{T_A}$ are given by $\lambda' = F' \lambda + (1 - F')/D$. Therefore, the trace norm can be calculated from $P(\tilde{\lambda})$ as

$$\begin{aligned} \|\hat{\rho}^{T_A}\|_1 &\approx 2F' D_A \left(\int_0^\infty \tilde{\lambda} P(\tilde{\lambda}) d\tilde{\lambda} \right. \\ &\quad \left. + \int_0^{(F' D_A)^{-1}} \left[(F' D_A)^{-1} - \tilde{\lambda} \right] P(\tilde{\lambda}) d\tilde{\lambda} \right). \end{aligned} \quad (9.16)$$

The first term evaluates to $\frac{64}{9\pi^2} F' D_A$, while the second term can be systematically computed by using the expansion $P(\tilde{\lambda}) \approx -\frac{2}{\pi^2} \ln |\tilde{\lambda}| + \frac{8 \ln 2 - 4}{\pi^2}$, valid for small $\tilde{\lambda}$, giving Eq. (9.13).

Importantly, if more terms are added in the expansion of Eq. (9.13) Lemma 1 is an exact estimator of the mixed state entanglement. Further, if D_A is large, as is the case for large system sizes, then the original simple proxy is restored as $f \sim F' \sim F$. We show the

efficacy of this estimator, with $f(F)$ evaluated to first order, in Fig. 9.4b.

Finally we prove an even more general set of upper and lower bounds. While not amenable to be computed experimentally, they allow us to argue our entanglement-proxy likely becomes exponentially tight as the system size increases.

Lemma 2. For any state pure state $|\psi\rangle$ and mixed state $\hat{\rho}$ with decomposition

$$\hat{\rho} = F|\psi\rangle\langle\psi| + (1 - F)\hat{\rho}_\perp, \quad (9.17)$$

where $\hat{\rho}_\perp$ is the (not necessarily positive semi-definite) remainder of the noisy state, we can bound the mixed state log negativity by

$$\begin{aligned} \mathcal{E}_N(\hat{\rho}) &\geq \log_2 \left(F2^{\mathcal{E}_N(|\psi\rangle)} - (1 - F)\|\hat{\rho}_\perp^{TA}\|_1 \right) \\ \mathcal{E}_N(\hat{\rho}) &\leq \log_2 \left(F2^{\mathcal{E}_N(|\psi\rangle)} + (1 - F)\|\hat{\rho}_\perp^{TA}\|_1 \right), \end{aligned} \quad (9.18)$$

where F is a lower bound for the fidelity $\langle\psi|\hat{\rho}|\psi\rangle$.

Proof. Unlike above, $\hat{\rho}_\perp$ does not necessarily commute with $|\psi\rangle\langle\psi|$.

To bound the negativity for such a state, we simply apply the triangle inequality on the trace norm $\|\cdot\|_1$, yielding Eqs. (9.18).

This bound is particularly useful in settings such as ours. Here, we can expect the noisy state to be weakly entangled, i.e. $\|\hat{\rho}_\perp^{TA}\|_1 = \mathcal{O}(1)$, while the pure state is highly entangled: $\|\psi\rangle\langle\psi|^{TA}\|_1 = \mathcal{O}(\exp(N))$, where N is the system size. Then the above bounds imply that our entanglement-proxy is in fact exponentially close to the true mixed state negativity: $\mathcal{E}_N(\hat{\rho}) = \mathcal{E}_P + \mathcal{O}(\exp(-N))$, as observed in the inset of Fig. 9.4b.

As a brief aside, we note an important point, arising from the fact that we have used either F_d as the fidelity for experiment, or the linear cross-entropy, F_{XEB} , as the fidelity for the literature examples. However, F_d and F_{XEB} are both affected by measurement error, which strictly speaking does not affect the mixed state entanglement. However, both F_d and F_{XEB} are expected to decrease due to measurement errors, and so this only serves to make our lower bound less tight. In principle the mixed state entanglement-proxy could be measurement-corrected to account for this effect, but we choose to include measurement errors as a more general quality-factor of the experiments.

9.3 ROBUSTNESS TO REALISTIC EXPERIMENTAL ERRORS

In the above section, we showed several proofs which let us argue our entanglement-proxy, \mathcal{E}_P is a tight lower bound for the mixed state entanglement under specific assumptions. Here, we relax these assumptions, but show that under realistic noise sources, maximum possible violations of our bound are at most $\mathcal{O}(1)$.

9.3.1 Error model simulations

To start, we reiterate and emphasize the results of Fig. 9.1a. There, we showed numerically that the entanglement-proxy appears to be a genuine lower bound of the mixed state entanglement either for a random unitary circuit (RUC) undergoing incoherent local noise, or for our experiment undergoing the full set of error sources (Fig. 8.12).

Now, our aim is to provide further analytical and numerical evidence for the robustness of the entanglement-proxy under: 1) global coherent errors, or 2) incoherent local errors.

9.3.2 Global coherent errors

In this section, we consider potential violations of our bound due to global coherent errors. To study this, we first prove a Theorem for a more general bound of the mixed state entanglement, based on studying the norm of the commutator between the mixed state $\hat{\rho}$ and pure state $|\psi\rangle$. This bound is equivalent to our original entanglement-proxy up to a correction term, which we then show analytically is small if the normalized entanglement of the target state is large. Finally, we present numerical support of this discussion.

Theorem 3 (Commutator based bound of the mixed state entanglement). *For any mixed state $\hat{\rho}$ and pure state $|\psi\rangle$ whose commutator has Frobenius norm $\mathcal{C} \equiv \|\llbracket \hat{\rho}, |\psi\rangle\langle\psi| \rrbracket\|_F$, the logarithmic negativities $\mathcal{E}_N(\hat{\rho})$ and $\mathcal{E}_N(|\psi\rangle)$ are related by*

$$\mathcal{E}_N(\hat{\rho}) \geq \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2 \log_2(1 - (C/F)\sqrt{(1-\alpha)/(2\alpha)}), \quad (9.19)$$

where $F \equiv \langle\psi|\hat{\rho}|\psi\rangle$ is the fidelity, $\alpha \equiv 2^{\mathcal{E}_N(|\psi\rangle)}/D_A \in [0, 1]$ is a normalized entanglement of $|\psi\rangle$, and D_A is the Hilbert space dimension of the subsystem A (assumed to be smaller than its complement).

Proof. We write $\hat{\rho}$ in a suitable basis as:

$$\hat{\rho} = \begin{pmatrix} F & B^\dagger \\ B & \hat{\rho}' \end{pmatrix}, \quad (9.20)$$

where the first row/column denotes the basis vector $|\psi\rangle$, and the second denotes the subspace perpendicular to $|\psi\rangle$. With P_\perp the projector onto this subspace, the matrix $\hat{\rho}' \equiv P_\perp \hat{\rho} P_\perp$ is positive semi-definite. The off-diagonal element $B \equiv P_\perp \hat{\rho} |\psi\rangle$ is a $(d-1) \times 1$ vector, which we take to be an unnormalized vector $|B\rangle$ in the full Hilbert space (with first entry equal to zero).

Using the variational definition of the trace-norm used in the proof of Theorem 2, we write

$$\|\hat{\rho}^{TA}\|_1 \geq \text{tr}(\hat{P}_\psi \hat{\rho}^{TA}) = \langle \Phi | \hat{\rho} | \Phi \rangle \quad (9.21)$$

$$= F |\langle \psi | \Phi \rangle|^2 + \langle \Phi | \hat{\rho}' | \Phi \rangle + \langle \Phi | \psi \rangle \langle B | \Phi \rangle + \langle \Phi | B \rangle \langle \psi | \Phi \rangle, \quad (9.22)$$

where $|\Phi\rangle$ is the *unnormalized* maximally entangled state in the full Hilbert space, in the Schmidt bases $\{|a_j\rangle\}$ and $\{|b_j\rangle\}$ of $|\psi\rangle$. We then bound the terms of this expression.

First, the Frobenius norm of the commutator $\mathcal{C} = \|\llbracket \hat{\rho}, |\psi\rangle\langle\psi| \rrbracket\|_F$ has the expression

$$\mathcal{C}^2 = 2\langle \psi | \hat{\rho}^2 | \psi \rangle - 2\langle \psi | \hat{\rho} | \psi \rangle^2 = 2\langle B | B \rangle. \quad (9.23)$$

Therefore, $\langle B | B \rangle \leq \mathcal{C}^2/2$. We then use the *Schur complement* condition for positive semi-definite matrices. With our decomposition, the fact that $\hat{\rho}$ is positive semi-definite implies that the Schur complement $\hat{\rho}' - F^{-1}|B\rangle\langle B|$ is positive semi-definite. We conclude that

$$\langle \Phi | B \rangle \langle B | \Phi \rangle \leq F \langle \Phi | \hat{\rho}' | \Phi \rangle. \quad (9.24)$$

Using the fact that, $|B\rangle \perp |\psi\rangle$ and that $|\Phi\rangle$ is an unnormalized state with norm D_A and overlap $|\langle \Phi | \psi \rangle|^2 \equiv \alpha D_A$, then accounting for the norms of $|\Phi\rangle$ and $|B\rangle$, the maximum possible value of $|\langle B | \Phi \rangle|$ is $|\langle B | \Phi \rangle| \leq \sqrt{(1-\alpha)D_A} \mathcal{C} / \sqrt{2}$. Eq. (9.24) also implies that $\langle \Phi | \hat{\rho}' | \Phi \rangle \geq (1-\alpha)D_A \mathcal{C}^2 / (2F)$.

Combining these ingredients with Eq. (9.22), we then have

$$\begin{aligned} \|\hat{\rho}^{TA}\|_1 &\geq \langle \Phi | \hat{\rho} | \Phi \rangle \\ &\geq F\alpha D_A + (1-\alpha)D_A \mathcal{C}^2 / (2F) - 2\sqrt{\alpha D_A^2 (1-\alpha) \mathcal{C}^2 / 2} \end{aligned}$$

$$\begin{aligned}
&= \left(\sqrt{F\alpha} - \sqrt{(1-\alpha)\mathcal{C}^2/(2F)} \right)^2 D_A \\
&= F\alpha D_A \left(1 - (\mathcal{C}/F)\sqrt{(1-\alpha)/(2\alpha)} \right)^2. \tag{9.25}
\end{aligned}$$

Taking the logarithm of both sides, we arrive at Theorem 3.

With this Theorem 3 in hand, to validate our entanglement proxy we need only bound the size of the quantity \mathcal{C}/F , and show it is $\mathcal{O}(1)$. Here, we show this is the case under the physically realistic scenario of the system undergoing coherent global errors, which are a dominant error source in our experiment (Fig. 8.12). First, we prove a useful Lemma rewriting the coefficient to the correction term, \mathcal{C}/F , in terms of only F .

Lemma 3. *For normally distributed global coherent quasistatic parameter errors, the coefficient to the correction term in Theorem 3, \mathcal{C}/F , can be written entirely in terms of the fidelity as*

$$\frac{\mathcal{C}}{F\sqrt{2}} = \sqrt{\frac{2}{\sqrt{3+2F^2-F^4}} - 1}. \tag{9.26}$$

Proof. To begin, consider evolution under a fixed parameter value, which results in the pure state $|\Psi(t, \theta)\rangle \equiv e^{-i(\hat{H}_0 + \theta\hat{V})t}|\Psi_0\rangle$. In Theorem 1, we found a Gaussian fidelity dependence between states corresponding to any two parameter values θ_1 and θ_2 :

$$F(\theta_1, \theta_2) \equiv |\langle \Psi(t, \theta_1) | \Psi(t, \theta_2) \rangle|^2 = \exp\left(-N\lambda t^2 \frac{(\theta_1 - \theta_2)^2}{2}\right), \tag{9.27}$$

For some constant λ [Eq. (8.5)]. Assuming a normal distribution $P(\theta)$ of parameter values θ with mean θ_0 and standard deviation σ , we obtain the mixed state $\hat{\rho}(t) = \int d\theta P(\theta) |\Psi(t, \theta)\rangle \langle \Psi(t, \theta)|$. In Eq. (8.19), we computed the fidelity of $\rho(t)$ to the pure state $|\Psi(t, \theta_0)\rangle$ and found it to be $F = 1/\sqrt{\Lambda + 1}$, where $\Lambda \equiv Nt^2\lambda\sigma^2$.

The Frobenius norm of the commutator, \mathcal{C} , then has expression

$$\mathcal{C}^2 = \|\hat{\rho}(t), |\Psi(t, \theta_0)\rangle \langle \Psi(t, \theta_0)|\|_F^2$$

$$\begin{aligned}
&= 2\langle \Psi(t, \theta_0) | \hat{\rho}(t)^2 | \Psi(t, \theta_0) \rangle - 2\langle \Psi(t, \theta_0) | \hat{\rho}(t) | \Psi(t, \theta_0) \rangle^2 \\
&= 2 \int d\theta_1 P(\theta_1) \int d\theta_2 P(\theta_2) \left(\langle \Psi(t, \theta_0) | \Psi(t, \theta_1) \rangle \right. \\
&\quad \langle \Psi(t, \theta_1) | \Psi(t, \theta_2) \rangle \\
&\quad \left. \langle \Psi(t, \theta_2) | \Psi(t, \theta_0) \rangle \right) - 2F^2 \\
&= \frac{4}{\sqrt{3\Lambda^2 + 8\Lambda + 4}} - \frac{2}{\Lambda + 1}. \tag{9.28}
\end{aligned}$$

In this computation, the unknown phases of the inner products $\langle \Psi(t, \theta_i) | \Psi(t, \theta_j) \rangle$ cancel, allowing the use of Theorem 1. Rearranging, it is straightforward to then arrive at Eq. (9.26).

We can then combine Theorem 3 and Lemma 3, to arrive at the simplified result

$$\mathcal{E}_N(\hat{\rho}) \geq \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2 \log_2 \left(1 - \sqrt{\left(\frac{2}{\sqrt{3 + 2F^2 - F^4}} - 1 \right) \frac{1 - \alpha}{\alpha}} \right) \tag{9.29}$$

$$\approx \begin{cases} \mathcal{E}_N(|\psi\rangle) + \log_2(F) & F \approx 1 \\ \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2 \log_2 \left(1 - 0.39 \sqrt{\frac{1 - \alpha}{\alpha}} \right) & F \ll 1 \end{cases} \tag{9.30}$$

which shows that our entanglement-proxy is robust in the presence of quasistatic parameter fluctuations, up to a $\mathcal{O}(1)$ deviation which decreases with the normalized entanglement $\alpha \equiv 2^{\mathcal{E}_N(|\psi\rangle)} / D_A$.

More concretely, let us assume a normalized entanglement of $\alpha = 0.8$ (as is approximately the expectation at $N = 60$ and $t = 14.3$ cycles, see Fig. 9.3). In that case, Eq. (9.30) reduces to

$$\mathcal{E}_N(\hat{\rho}) \gtrsim \mathcal{E}_P(\hat{\rho}) - 0.6. \tag{9.31}$$

In other words, we expect the maximum possible violation of our bound is at most around half an ebit of entanglement. We emphasize that this does not mean the bound is in fact violated by this amount – only that it *potentially could be*.

In order to validate that the coefficient to the correction term is small, in Fig. 9.5a we show $\mathcal{C}/(F\sqrt{2})$ as a function of F for various system sizes, calculated with our full Rydberg error model (which includes both Markovian and non-Markovian terms, and both global and local errors). Despite the error model including more noise terms than just global parameter fluctuations, we see curves as a function of N increase but appear to converge

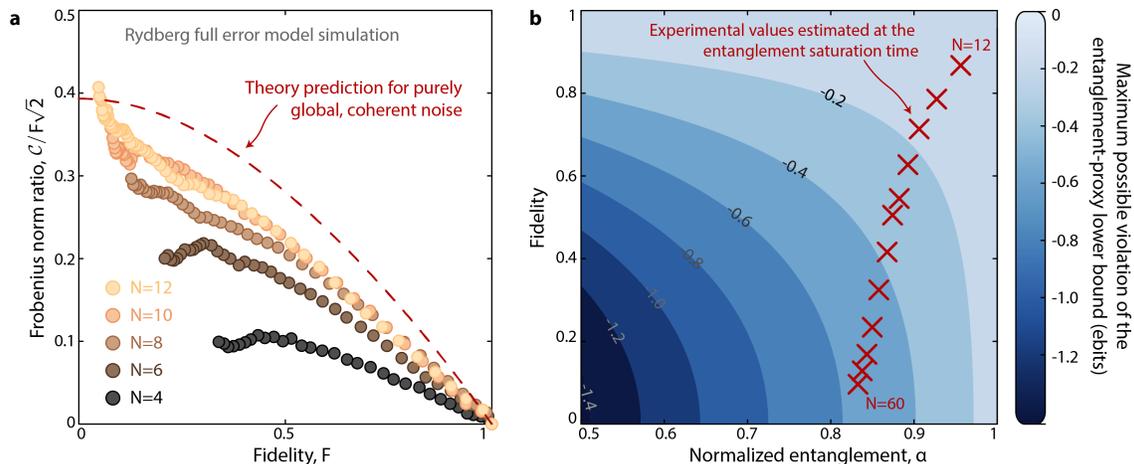


Figure 9.5: Bounding the entanglement-proxy in the presence of errors. **a.** We consider full error model simulations of the Rydberg dynamics, i.e. simultaneously under both global and local errors which each can be either Markovian or non-Markovian. We study the ratio of the commutator norm to the fidelity versus fidelity, truncating to only show data for which $F > 1/D_A$, where D_A is the half chain Hilbert space dimension. The data appears to converge as a function of system size, approaching the theory prediction obtained for the case of purely global, coherent Hamiltonian fluctuations, see Eq. (9.29). **b.** Here we plot the expected maximum violation of our entanglement-proxy lower bound as a function of the fidelity and the normalized entanglement; more specifically, we plot the correction term from Eq. (9.29). We find that in the reasonable parameter regime shown, the violation is always $\lesssim 1$ ebit. From analysis shown in Fig. 9.3, we plot the expected experimental values from $N = 12$ to 60 , and find the maximum violation appears to saturate around ≈ 0.6 ebits. Together these results imply that our entanglement-proxy is robust up to $\mathcal{O}(1)$ corrections in the face of realistic noise.

near the prediction of Eq. (9.26). Then, assuming the prediction of Eq. (9.26), we show the maximum possible violation of the lower bound provided by \mathcal{E}_P (i.e. the correction term in Eq. (9.29)) as a function of normalized entanglement and fidelity. We further include estimated experimental values up to $N = 60$ (Fig. 9.5b). We find the experimental values appears to saturate around a maximum possible violation of ≈ 0.6 ebits.

9.3.3 Systematic miscalibration errors

Next, we comment on the case where the target state $|\psi\rangle$ is set incorrectly, for instance because of a miscalibration of the Hamiltonian parameters. While this type of error can lead to a violation of our lower bound, the violation is small and the period for which it occurs is ephemeral, because the incorrect choice of $|\psi\rangle$ will lead to a fidelity decay which is stronger than any typical increase in pure state entanglement. To showcase this, we consider

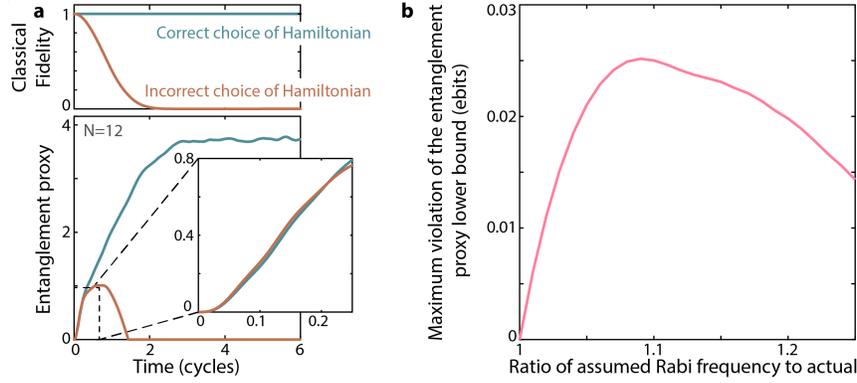


Figure 9.6: **Sensitivity to incorrect pure state.** **a.** We study a possible error source in our entanglement proxy, that the pure state negativity could be assumed incorrectly, for instance if the wrong Hamiltonian parameters are chosen in simulation. We consider error-free quench evolution with $\Omega/2\pi = 6.9$ MHz, leading to normal linear growth and then saturation of pure state negativity (which is equal to the entanglement proxy, as $F = 1$). If we calculated the entanglement proxy assuming the Rabi frequency was instead $\Omega/2\pi = 7.2$ MHz, then at very early times the calculated entanglement proxy can actually slightly exceed the true entanglement, showcasing an adversarial example where our entanglement proxy fails to be a lower bound (inset). However, loss of fidelity due to the Hamiltonian parameter mismatch will eventually dominate (top), restoring the validity of the lower bound for large system sizes or late evolution times. **b.** Even for a wide range of possible Rabi frequency miscalibrations (represented by the ratio of the miscalibrated Rabi frequency to the actual Rabi frequency), the maximum possible violation of the proxy lower bound is only ~ 0.02 ebits.

quench dynamics resulting in linear growth of negativity followed by saturation (Fig. 9.6). We study the two entanglement regimes (saturated or growing linearly) separately.

First, we consider the case where the entanglement has saturated. We imagine the true evolution was performed with a Rabi frequency Ω_0 , resulting in a state $|\psi_0\rangle$ at time t . However, we consider the situation where the parameters in the Hamiltonian have been mischosen, for instance with Rabi frequency $\Omega_1 = \Omega_0 + \omega$. From Theorem 1 we know the resulting fidelity overlap between the two states is given by

$$F = |\langle\psi_0|\psi_1\rangle|^2 = 2^{-N\lambda\omega^2 t^2/2}. \quad (9.32)$$

Because we are studying the entanglement saturated regime, the pure state negativities of $|\psi_0\rangle$ and $|\psi_1\rangle$ are approximately equal, but the mixed state negativity will now be calculated as

$$\mathcal{E}_P(|\psi_1\rangle) = \mathcal{E}_N(|\psi_1\rangle) + \log_2(F) \quad (9.33)$$

$$\approx \mathcal{E}_P(|\psi_0\rangle) - N\lambda\omega^2 t^2/2. \quad (9.34)$$

Therefore, any change in the assumed Rabi frequency will lead to a quadratic loss in the mixed state entanglement proxy, and the proxy will remain a lower bound.

The situation is largely the same if we now work in the linear growth regime, where we assume the pure state entanglement grows as $\mathcal{E}(|\psi_0\rangle) = \alpha\Omega_0 t$, and $\mathcal{E}(|\psi_1\rangle) = \alpha\Omega_1 t$, where α is the entanglement velocity related to Lieb-Robinson bounds. In this case, the entanglement proxy can be written as

$$\mathcal{E}_P(|\psi_1\rangle) = \mathcal{E}_N(|\psi_1\rangle) + \log_2(F) \quad (9.35)$$

$$= \alpha\Omega_1 t - N\lambda\omega^2 t^2/2 \quad (9.36)$$

$$= \mathcal{E}_P(|\psi_0\rangle) + \alpha\omega t - N\lambda\omega^2 t^2/2 \quad (9.37)$$

$$= \mathcal{E}_P(|\psi_0\rangle) + \omega t(\alpha - N\lambda\omega t/2). \quad (9.38)$$

Thus, as long as $N\lambda\omega t/2 > \alpha$, the erroneous choice of target state will still result in a lower overall mixed state entanglement proxy, thus maintaining the lower bound. This will then become increasingly tight for larger system sizes and later evolution times. In Fig. 9.6 we show simulations demonstrating this, indicating that Hamiltonian parameter miscalibration should not substantially affect the veracity of our mixed state entanglement lower bound past very early times.

9.3.4 Incoherent local errors

We now turn to prove our entanglement proxy is a valid lower bound for the case of random unitary circuits (RUCs) undergoing realistic noise, namely local depolarization error. To do so, we will introduce a new theorem, which bounds the violation in our entanglement proxy by a value related to the fidelity overlap with eigenstates of $\hat{\rho}$, and then will show that this correction term is small.

We proceed first by proving a more general Lemma which will assist us; the proof has similarities to the proof of Theorem 3, but we will present it again here for completeness.

Lemma 4. *For any mixed state $\hat{\rho}$ and pure state $|\psi\rangle$*

$$\mathcal{E}_N(\hat{\rho}) \geq \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2\log_2(1 - \sqrt{(1-F)(1-\alpha)/(F\alpha)}), \quad (9.39)$$

where $F \equiv \langle\psi|\hat{\rho}|\psi\rangle$ is their fidelity, $\alpha \equiv 2^{\mathcal{E}_N(|\psi\rangle)}/D_A \in [0, 1]$ is a normalized entanglement of $|\psi\rangle$, and D_A is the Hilbert space dimension of the subsystem A (assumed to be smaller than its complement).

Proof. In deriving Theorem 2, we utilized the fact that when $|\psi\rangle$ is an eigenstate of $\hat{\rho}$, the remainder state $\hat{\rho}_\perp$ is positive semi-definite, which allowed us to bound $\langle\Phi|\hat{\rho}_\perp|\Phi\rangle \geq 0$. To show Lemma 4, we use constraints that arise from the fact that $\hat{\rho}$ is positive semidefinite. We write $\hat{\rho}$ in a suitable basis as:

$$\hat{\rho} = \begin{pmatrix} F & B^\dagger \\ B & \hat{\rho}' \end{pmatrix}, \quad (9.40)$$

where the first row/column denotes the basis vector $|\psi\rangle$, and the second denotes the subspace perpendicular to $|\psi\rangle$. With P_\perp the projector onto this subspace, the matrix $\hat{\rho}' \equiv P_\perp \hat{\rho} P_\perp$ is positive semi-definite. The off-diagonal element $B \equiv P_\perp \hat{\rho} |\psi\rangle$ is a $(d-1) \times 1$ vector, which we take to be an unnormalized vector $|B\rangle$ in the full Hilbert space (with first entry equal to zero).

Using Eq. (9.7), the negativity can be obtained from

$$\|\hat{\rho}^{TA}\|_1 \geq \langle\Phi|\hat{\rho}|\Phi\rangle = F |\langle\psi|\Phi\rangle|^2 + \langle\Phi|\hat{\rho}'|\Phi\rangle + \langle\Phi|\psi\rangle\langle B|\Phi\rangle + \langle\Phi|B\rangle\langle\psi|\Phi\rangle. \quad (9.41)$$

Only the last two terms above may be negative. We bound their magnitude using the *Schur complement* condition for positive semi-definite matrices. With our decomposition, the fact that ρ is positive semi-definite implies that the Schur complement $\hat{\rho}' - F^{-1}|B\rangle\langle B|$ is positive semi-definite. We conclude that

$$\langle\Phi|B\rangle\langle B|\Phi\rangle \leq F \langle\Phi|\hat{\rho}'|\Phi\rangle, \quad (9.42)$$

which allows us to further conclude that

$$\langle\Phi|\hat{\rho}|\Phi\rangle \geq (\sqrt{F} |\langle\Phi|\psi\rangle| - \sqrt{\langle\Phi|\hat{\rho}'|\Phi\rangle})^2. \quad (9.43)$$

We then use the fact that $|\Phi\rangle/\sqrt{D_A}$ is a normalized state. Since $|\langle\Phi|\psi\rangle|^2 = \alpha D_A = \|\psi\rangle\langle\psi|^{TA}\|_1$, $\text{tr}\hat{\rho}' = 1 - F$ and $\langle\psi|\hat{\rho}'|\psi\rangle = 0$, we conclude that

$$\langle\Phi|\hat{\rho}'|\Phi\rangle \leq (1 - F)(1 - \alpha)D_A. \quad (9.44)$$

Then when $F \geq 1 - \alpha$, we obtain the claimed Eq. (9.39).

We then simply generalize Lemma 4 to the case where $|\psi\rangle$ has fidelity overlap with a particular eigenstate of $\hat{\rho}$.

Theorem 4 (Eigenstate-fidelity based bound). *For any mixed state $\hat{\rho}$ and pure state $|\psi\rangle$ which has a fidelity $f \equiv |\langle\psi|\lambda\rangle|^2$ with an eigenstate $|\lambda\rangle$ of $\hat{\rho}$, the logarithmic negativities $\mathcal{E}_N(\hat{\rho})$ and $\mathcal{E}_N(|\psi\rangle)$ are related by*

$$\mathcal{E}_N(\hat{\rho}) \geq \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2 \log_2(1 - \sqrt{(1-f)(1-\alpha)/(f\alpha)}), \quad (9.45)$$

where $F \equiv \langle\psi|\hat{\rho}|\psi\rangle$ is the fidelity, $\alpha \equiv 2^{\mathcal{E}_N(|\psi\rangle)}/D_A \in [0, 1]$ is a normalized entanglement of $|\psi\rangle$, and D_A is the Hilbert space dimension of the subsystem A (assumed to be smaller than its complement).

Proof. We first write $\hat{\rho} = F_t|v\rangle\langle v| + (1 - F_t)\hat{\rho}_\perp$, where $|v\rangle$ is an eigenstate of $\hat{\rho}$ with eigenvalue F_t , and $\hat{\rho}_\perp$ is the state projected onto the orthogonal subspace to $|v\rangle$. We write the state $|\psi\rangle$ that we benchmark against $\hat{\rho}$ in terms of $|v\rangle$ as $|\psi\rangle = \sqrt{f}|v\rangle + \sqrt{1-f}|\psi_\perp\rangle$. Then writing $\hat{\rho}$ in the same basis as Eq. (9.20) gives

$$\hat{\rho} = \begin{pmatrix} F_t f & F_t \sqrt{f(1-f)}|\psi\rangle\langle\psi_\perp| \\ F_t \sqrt{f(1-f)}|\psi_\perp\rangle\langle\psi| & \hat{\rho}_\perp + F_t(1-f)|\psi_\perp\rangle\langle\psi_\perp| \end{pmatrix}. \quad (9.46)$$

As with the above, the only negative contribution to $\langle\Phi|\hat{\rho}|\Phi\rangle$ arises from the off-diagonal terms of Eq. (9.46). This allows us to improve our bound to the claimed Theorem 4.

We now show that under incoherent local errors, the correction term presented in Theorem 4 is small. We study the case of random unitary circuits (RUCs) in the presence of local depolarization in order to treat the problem analytically.

We study the fidelity and purity and bound the fidelity $f \equiv |\langle\psi(t)|v\rangle|^2$ between the ideal evolved state $|\psi(t)\rangle$ with the largest eigenstate $|v\rangle$ of the noisy evolved state $\hat{\rho}(t)$, by computing both the purity and fidelity of $|\psi(t)\rangle$ and $\hat{\rho}(t)$. The relationship between these quantities in fact bounds the fidelity of $|\psi(t)\rangle$ to the largest eigenstate of $\hat{\rho}$, allowing us to apply Theorem 4 to evaluate the robustness of our negativity proxy.

Physically, our result means that the state $\hat{\rho}$ is a mixture of a particular pure state $|v_1\rangle$ with large population $F = \lambda_1$ and many other orthogonal states with smaller populations. For a circuit with spacetime volume $V = Nt$, the largest eigenvalue decays exponentially $\lambda_1 \approx (1-p)^V$, yet remains larger than all other eigenvalues. The other states correspond to trajectories with errors. The average number of errors is $k = pV$, and there are roughly

$\sigma_1, \sigma_2, \sigma_3$	$W_P(\sigma_1, \sigma_2, \sigma_3; p)$	$W_F(\sigma_1, \sigma_2, \sigma_3; p)$
+, +, +	1	1
+, +, -	0	0
+, -, +	$\frac{d}{d^2+1} + \frac{2d}{d^4-1}p + \frac{4d}{3(d^4-1)}p^2$	$\frac{d}{d^2+1} + \frac{d}{d^4-1}p$
+, -, -	$\frac{d}{d^2+1} + \frac{-2d^3}{d^4-1}p + \frac{4d^3}{3(d^4-1)}p^2$	$\frac{d}{d^2+1} + \frac{-d^3}{d^4-1}p$
-, +, +	$\frac{d}{d^2+1} + \frac{2d}{d^4-1}p + \frac{4d}{3(d^4-1)}p^2$	$\frac{d}{d^2+1} + \frac{d}{d^4-1}p$
-, +, -	$\frac{d}{d^2+1} + \frac{-2d^3}{d^4-1}p + \frac{4d^3}{3(d^4-1)}p^2$	$\frac{d}{d^2+1} + \frac{-d^3}{d^4-1}p$
-, -, +	$\frac{4d^2}{d^4-1}p + \frac{-20d^2}{3(d^4-1)}p^2 + \frac{16d^2}{3(d^4-1)}p^3 + \frac{-16d^2}{9(d^4-1)}p^4$	$\frac{2d^2}{d^4-1}p + \frac{-d^2}{d^4-1}p^2$
-, -, -	$1 + \frac{-4d^4}{d^4-1}p + \frac{20d^4}{3(d^4-1)}p^2 + \frac{-16d^4}{3(d^4-1)}p^3 + \frac{16d^4}{9(d^4-1)}p^4$	$1 + \frac{-2d^4}{d^4-1}p + \frac{d^4}{d^4-1}p^2$

Table 9.2: Ising Boltzmann weights of local downward-facing triangle configurations for the effective statistical mechanical models. Their partition functions equal the purity P [Eq. (9.47)] and fidelity F [Eq. (9.48)] of random unitary circuits (RUCs) with local depolarization. The geometric arrangement of $\sigma_1, \sigma_2, \sigma_3$ is indicated in Fig. 9.7c. In this work, we consider RUCs of qubits, with local Hilbert space dimension $d = 2$, and only consider the $\mathcal{O}(1)$ and $\mathcal{O}(p)$ terms. In this table, however, we state the full expressions in terms of general d for future reference.

$\frac{V!}{k!(V-k)!}$ different spacetime locations of k errors. When the number of trajectories are much less than the Hilbert space dimension D , we expect the resultant wavefunctions resultant to be orthogonal to one another. This is true if the k errors are sparsely spread over the system in spacetime (such that their separations are sufficiently large) and the effects of errors are scrambled quickly. When p is sufficiently small and when the quantum dynamics is sufficiently chaotic, these assumptions are expected to hold.

We proceed to justify this intuition by averaging over random circuits. At every time step, we apply random unitary two-qubit gates in a brickwork circuit geometry, and subject it to local depolarization, which locally depolarizes the state ρ at every site with probability p , equivalent to applying the local channel $\rho \mapsto (1-p)\rho + p(I_j/d) \otimes \text{tr}_j(\rho)$ independently to every site j , where $d = 2$ is the local Hilbert space dimension (Fig. 9.7a). For simplicity, we also take the initial state $|\psi_0\rangle$ to be the product state $|0\rangle^{\otimes N}$. After averaging over RUCs, the purity $P \equiv \text{tr}[\rho(t)^2]$ and fidelity $F \equiv \langle \psi(t) | \rho(t) | \psi(t) \rangle$ can be expressed as the partition functions of two different Ising ferromagnets, with the Ising spins $\sigma = \pm 1$ respectively denoting the local identity and swap permutations (Fig. 9.7b).

The derivation of the mapping between RUCs and classical spin models is described in detail in many works including Refs. [420, 421]. Here we simply quote the result:

$$P = \sum_{\{\sigma\}} \prod_{\nabla} W_P(\sigma_1, \sigma_2, \sigma_3; p), \quad (9.47)$$

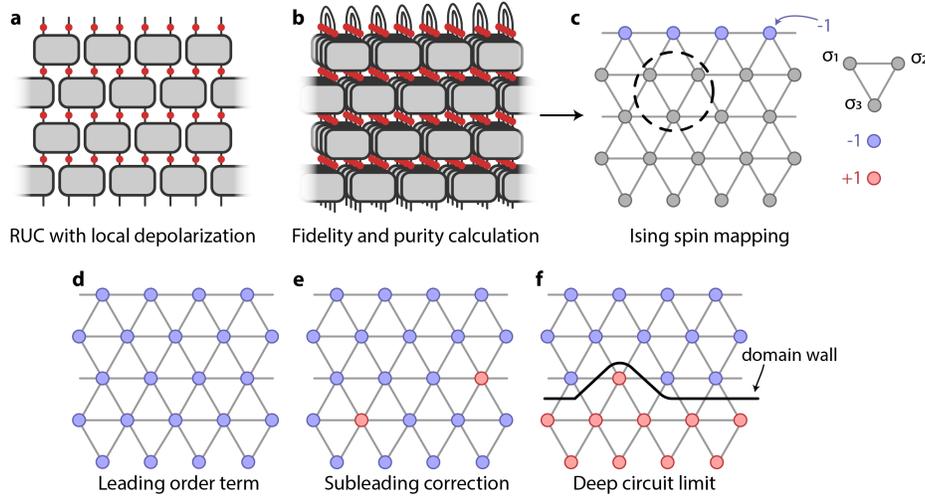


Figure 9.7: Purity and fidelity in random unitary circuits (RUCs) with local depolarization. **a.** We consider a brickwork RUC, where local depolarization with strength p is applied after every layer and at every site (red dots). **b.** Typical values of purity or fidelity can be computed by random circuit averages. In both quantities, swap boundary conditions are applied on four copies of the circuit, representing two copies of the states. For purity, dephasing is applied to both copies (depicted here), while for fidelity, dephasing is only applied to one copy. **c.** After averaging over random circuits, the purity and fidelity can be expressed as the partition functions of effective statistical mechanical Ising models, with three-spin Boltzmann weights on every downward-facing triangle (listed in Table 9.2), the top boundary fixed to be ‘-’ (blue circles) and the bottom boundary free. **d.** The dominant Ising spin configuration for both quantities is where all Ising spins are in the ‘-’ (swap) state (blue circles). **e.** Subleading corrections arise from configurations with sparse, isolated “bubbles” of ‘+’ Ising spins (red circles). **f.** For sufficiently deep circuits, the dominant configurations are those with a boundary domain of ‘-’ spins and a bulk domain of ‘+’ spins. However, this is only relevant in a regime where the fidelity is trivially small.

$$F = \sum_{\{\sigma\}} \prod_{\nabla} W_F(\sigma_1, \sigma_2, \sigma_3; p), \quad (9.48)$$

where the summation is over all possible Ising configurations of classical spin variables $\sigma = \pm 1$ and the product is over all downward-facing triangles in the resulting triangular lattice (Fig. 9.7c). $W_{P/F}(\sigma_1, \sigma_2, \sigma_3; p)$ act as local “Boltzmann weights,” parameterized by the error probability p . Their values are enumerated in Table 9.2. The Ising spins at the bottom are determined by the initial condition and are effectively unconstrained (open boundary condition), while the purity and fidelity expressions result in an additional layer of Ising spins at the top, which are pinned to be the swap permutation ‘-’.

When $p = 0$, notice that $W_{P/F}(-, -, +) = 0$. Combined with the fact that all top boundary spins have $\sigma = -$, the only non-vanishing contribution to P or F is the one where all $\sigma = -$

(Fig. 9.7d). In this case, $P = \prod_{\square} W_P(-, -, -) = 1$ and similarly $F = 1$. This is expected since in the absence of errors, ρ is a pure state with perfect fidelity.

Now consider small, nonzero p . P and F have contributions from configurations where some of the spin variables are flipped to the ‘+’ state, with weight $\mathcal{O}(p)$ (Fig. 9.7e). We call such spin flips “bubbles,” because such bubbles are penalized from growing by an effective line-tension term, which weights larger bubbles by a factor of $(2/5)^l$, where l is the length of the boundary of the bubble. These small bubbles will effectively have renormalized weights, but will remain $\mathcal{O}(p)$ and whose precise values are not consequential to our conclusions. Therefore, in this regime, the number of such spin flips will be rare and sparse. To this end, we approximate P and F by simply summing over the contributions from sparse, disjoint bubbles. However, this approximation breaks down at larger values of p , discussed below. Within this approximation,

$$P \approx \left(1 - \frac{4d^2}{d^4 - 1}p + \mathcal{O}(p^2)\right)^V \exp(A_P V) \equiv \exp(-B_P V), \quad (9.49)$$

$$F \approx \left(1 - \frac{2d^2}{d^4 - 1}p + \mathcal{O}(p^2)\right)^V \exp(A_F V) \equiv \exp(-B_F V), \quad (9.50)$$

where V is the spacetime volume (number of Ising spins) and A_F and A_P are the renormalized relative weights of a single bubble.

We are interested in the ratio $P/F^2 \approx e^{(2B_F - B_P)V} \geq 1$. When $p \ll 1$,

$$\begin{aligned} 2B_F - B_P &= -2A_F + A_P - 2 \ln \left(1 - \frac{2d^2}{d^4 - 1}p + \mathcal{O}(p^2)\right) + \ln \left(1 - \frac{4d^2}{d^4 - 1}p + \mathcal{O}(p^2)\right) \\ &= A_P - 2A_F + \mathcal{O}(p^2). \end{aligned} \quad (9.51)$$

Because each bubble already has $\mathcal{O}(p)$ weight, the leading order contribution to $A_P - 2A_F$ is determined by that of $W_P(-, -, +)$ and $W_F(-, -, +)$, which have ratio precisely 2. Therefore, we conclude that all $\mathcal{O}(p)$ terms vanish and

$$P/F^2 = e^{Cp^2V}, \quad (9.52)$$

for some $\mathcal{O}(1)$ constant C .

In order to translate this into a bound on the eigenstate-fidelity f , we denote the fidelities between $|\psi\rangle$ and the eigenstates $|v_i\rangle$ of $\hat{\rho}$ as $\mu_i \equiv |\langle v_i | \psi \rangle|^2$, with $f \equiv \mu_1$. With λ_i the eigenvalues of $\hat{\rho}$, we have $F = \sum_i \lambda_i \mu_i$ and $P = \sum_i \lambda_i^2$. The Cauchy-Schwarz inequality implies that

$$\mu_1^2 + (1 - \mu_1)^2 \geq \sum_i \mu_i^2 \geq F^2/P \Rightarrow 1 - \mu_1 \leq \frac{P/F^2 - 1}{2P/F^2}. \quad (9.53)$$

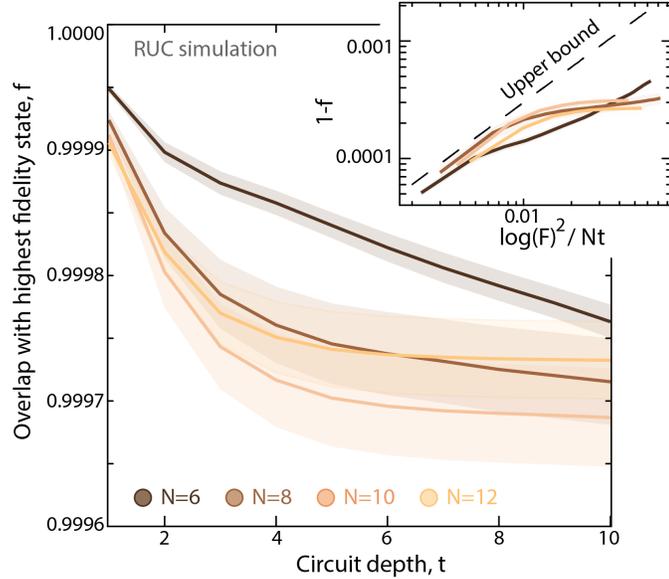


Figure 9.8: **Bounding the entanglement-proxy in the presence of errors.** We consider the case of incoherent local errors, specifically for random unitary circuits undergoing local depolarization. The overlap, f , between the target pure state and the highest fidelity state to $\hat{\rho}$ remains near unity. Further, $1 - f$ is linearly bounded by the quantity $\log(F)^2/Nt$ with a small proportionality constant 0.03 (inset), consistent with our RUC analysis leading up to Eq. (9.54), and giving the overall bound Eq. (9.55).

Using Eq. (9.52), we obtain

$$1 - f \equiv 1 - \mu_1 \leq (1 - \exp(-Cp^2V))/2 = \mathcal{O}(p^2V). \quad (9.54)$$

Using Theorem 4, we obtain

$$\mathcal{E}_N(\hat{\rho}) \gtrsim \mathcal{E}_N(|\psi\rangle) + \log_2(F) + 2 \log_2 \left(1 - \sqrt{\frac{1-\alpha}{\alpha}} C_2 \frac{\ln(F)^2}{V} \right), \quad (9.55)$$

for another $\mathcal{O}(1)$ constant C_2 , indicating that the entanglement-proxy E_P is valid up to a small correction which is negligible as long as $F \gg \exp(-\sqrt{V}) = \exp(-\sqrt{Nt}) \sim \exp(-N)$, i.e. F is not trivially small. We give numerical support for this in Fig. 9.8.

Finally, we remark that additional configurations may come to dominate the partition function sums Eqs. (9.47), (9.48) when the circuit is sufficiently deep. This is easiest to understand in the statistical mechanical picture. For both the purity and fidelity, the bulk has a small preference for the ‘+’ state, because depolarization breaks the Ising symmetry and favors the identity element. Meanwhile, the boundary condition (associated with evaluating purity or fidelity) pins the boundary to the ‘-’ state, acting as a strong boundary

field. When the circuit is shallow, the boundary term dominates, giving rise to the global ‘-’ domain assumed above. However, when the circuit is sufficiently deep, it becomes energetically favorable for a large domain wall to form, dividing the spins into a ‘-’ domain near the boundary, and a ‘+’ domain in the bulk (Fig. 9.7f), at the expense of a domain wall with free-energy scaling with the spatial size. In our above calculations, we neglect this phenomenon because it is only relevant when the fidelity is trivially small, $\mathcal{O}(D^{-1})$, where D is the total Hilbert space dimension. In this regime, the purity and fidelity decouple from the bulk dynamics and assume this constant value which arises from the boundary Ising ‘-’ domain.

9.4 ENTANGLEMENT OF STATE PREPARATION

The mixed state entanglement proxy we have introduced (Eq. 9.1) is a general lower bound in generic situations for any target pure state of interest. We have demonstrated its applicability in quantum quench dynamics, and we anticipate its direct applicability in a much broader range of quantum simulation experiments. To exemplify this, we consider in this section the estimation of mixed state entanglement during quasi-adiabatic state preparation and at quantum critical points.

To estimate the mixed state entanglement with our proxy, we require knowledge of the entanglement of the pure target state and the fidelity of experimental state to the pure target state. At quantum critical points, the scaling of entanglement is well studied [422], and the pure state entanglement can be extrapolated from exact simulations at small system size.

It remains then to estimate the fidelity, for which we identify two possible methods (Fig. 9.9a):

Method A) For particular states of interest, F_d may not work directly. However, F_d can still be applied by following the quasi-adiabatic sweep preparation with a short, ergodic quench [29, 30]. Importantly, F_d is accurate for short quenches, meaning there is little fidelity loss in this process, and that tensor network methods can be used to generate exact classical simulation references even for large system sizes where critical states have area law entanglement with potentially logarithmic corrections [422]. Further, the performance of this method could be improved by calibrating the fidelity loss due to the quench via measurements at multiple times.

Method B) Continue the adiabatic sweep into a phase with a ground state for which the fidelity is easily measured (such as a product state). The fidelity of preparing the critical state can then be approximated as the fidelity at the end of the sweep.

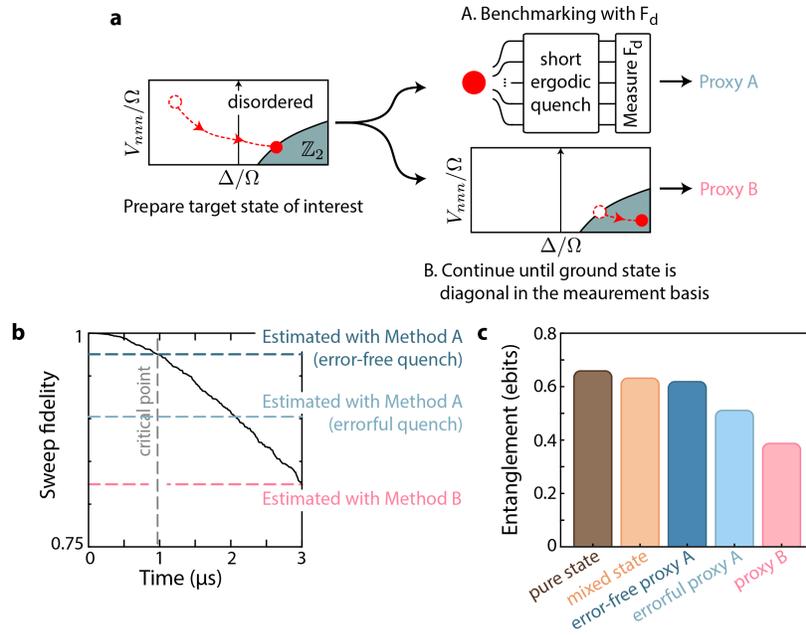


Figure 9.9: Estimating mixed state entanglement of adiabatic state preparation. **a.** The mixed state entanglement proxy can be applied to the case of adiabatic state preparation. For instance, a quasi-adiabatic sweep may be performed to prepare the critical state at a quantum phase transition. To estimate the fidelity, there are two possible methods: A) use the F_d formula presented here (and as further discussed in this context of *target state benchmarking* in Ref. [29]), or B) continue the sweep into a classical phase for which the ground state fidelity can be directly measured. **b, c.** We use our error model to simulate a sweep to the quantum phase transition between the disordered and \mathbb{Z}_2 phases of the ground state Rydberg phase diagram in 1D. For method A, after reaching the critical point we performed a disordered quench for 4 cycles and calculate F_d to estimate the fidelity. For method B, we use the \mathbb{Z}_2 population at the end of the entire sweep to estimate the fidelity. While both options work, for the particular sweep simulated here, method A yields a tighter lower bound on the actual mixed state negativity at the critical point, particularly if the measurement quench is error-free.

We demonstrate both fidelity estimation methods using our standard error model (with no preparation or measurement errors) to simulate a quasi-adiabatic sweep to prepare the ground state of the quantum phase transition between the disordered and classical \mathbb{Z}_2 ordered phases in the 1D Rydberg phase diagram (Fig. 9.9b) with $N = 11$ atoms. The sweep in question has a tangent detuning profile with maximum detuning range of $\pm 2\pi \times 40$ MHz, and Rabi frequency of $2\pi \times 5$ MHz. The sweep duration is $3 \mu\text{s}$.

For Method A, we simulate the preparation of the target state at the critical point. We then use the error model to simulate a quench for 4 cycles with a time-independent Hamiltonian with $\Omega = 2\pi \times 6.9$ MHz, and with $2\pi \times \pm 2$ MHz random on-site disorder drawn from a

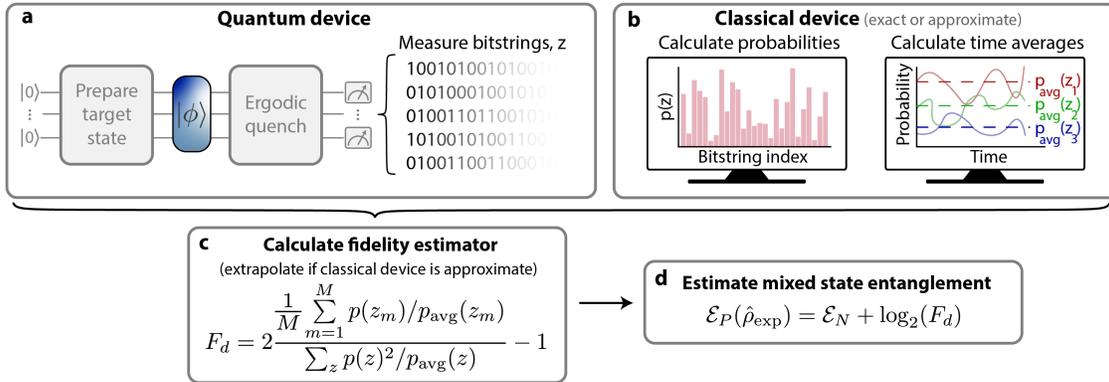


Figure 9.10: Schematic for mixed state entanglement estimation of a target quantum state. **a.** The quantum device prepares a target state [29] of interest, $|\psi\rangle$, before performing an ergodic quench, and measuring a set of resulting bitstrings. The experimental fidelity drops both during state preparation and the quench due to errors (not pictured). **b.** A classical computer then performs noise-free simulations of the entire dynamics (either exactly or approximately), calculates both the bitstring probabilities and the time-averaged bitstring probabilities, p_{avg} , and further estimates the second moment of the bitstring probability distribution. **c.** The fidelity is then estimated via a cross-entropy type quantity [29, 30, 375], F_d . If the classical simulation were approximate, F_d is predicted either via Monte Carlo inference or direct extrapolation. **d.** The negativity of the target state is estimated or assumed, and the experimental mixed state entanglement-proxy is calculated.

uniform distribution. Such a technique in a similar parameter regime has been well studied theoretically in our previous works [29, 30]. For Method B, we measure the population of the classical \mathbb{Z}_2 ordered state ($|10101 \dots 0101\rangle$) at the end of the entire quench. Similar techniques can also be used for estimating fidelities of two-dimensional arrays [42, 43]. For both methods the mixed state entanglement proxy remains a lower bound to the actual mixed state negativity at the quantum phase transition, with Method A providing a tighter bound. For Method A we also provide the value assuming the measurement quench is error-free (or where the fidelity loss due to the quench is calibrated away), which closely estimates the actual mixed state negativity (Fig. 9.9c). We summarize the entirety of Method A in Fig. 9.10.

CHAPTER 10

The classical cost of quantum simulation

We finally ask: which device, quantum or classical, has a higher fidelity of reproducing a high-entanglement pure target state of interest? Equivalently, in term of fidelity, what are the minimum classical resources required for a classical computer to outperform the quantum device?

10.1 OVERVIEW

To answer this question, we compare the fidelity of the experiment against that of the MPS with varying bond dimension. We define the critical bond dimension for a given system size, χ^* , as the minimum bond dimension for which the classical fidelity always exceeds the estimated experimental fidelity. Importantly, this controls the costs of classical simulation – for instance, MPS simulation time scales as $\mathcal{O}(N\chi^3)$. We find χ^* continually increases as a function of system size (Fig. 10.1a), reaching a maximum value of $\chi^* = 3400$ for $N = 60$ (Fig. 10.10), and apparently continuing to increase beyond that point.

In performing this study, we used our new Lightcone-MPS algorithm, but considered several alternative approximate classical algorithms, including path integral [162], matrix product operator [424], time-dependent variational principle [425], Schrieffer-Wolff transformation [97], and neural net [426] approaches (see Section 10.2); however, we found the equivalent classical cost of these methods quickly became infeasible, typically well before $N = 60$. As an example, we show χ^* for a more conventional MPS approach using time-evolving block decimation [423] (Fig. 10.1a).

All calculations used a single 16-core node of the Caltech central computing cluster (see Section 10.2). On this machine, we estimate that running the Lightcone-MPS simulation

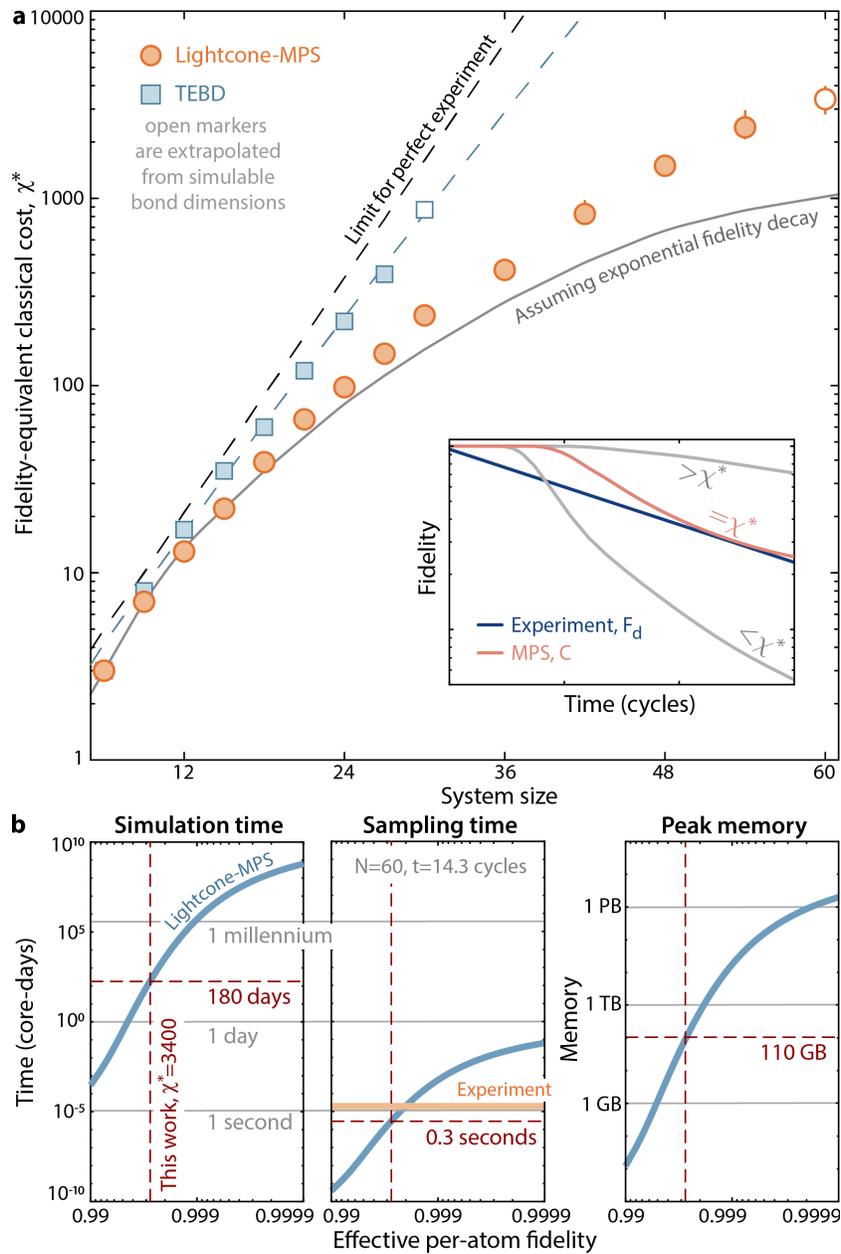


Figure 10.1: Classical cost to simulate the experiment. **a.** The equivalent classical cost of the experiment, as quantified by the minimum bond dimension, χ^* , for the classical simulation to maintain a higher fidelity than the experiment across all times, i.e. for $C > F_d$ (inset). We consider several classical algorithms (for example, time-evolving block decimation, TEBD [423]), all of which become impractical at moderate system sizes. This necessitates the introduction of our Lightcone-MPS algorithm (Section 10.2), which reaches a maximum value of $\chi^* = 3400$ for $N = 60$. **b.** Predicted MPS costs (simulation time, sampling time, and peak memory usage) to operate at χ^* as a function of the experimental per-atom fidelity (see text). Times are representative of a single 16-core node on the Caltech cluster.

for $N = 60$ and $\chi^* = 3400$ would entail a peak memory usage of ~ 110 GB (scaling as $\mathcal{O}(N\chi^2)$), and would take ~ 11.3 days, or $11.3 \times 16 \approx 180$ core-days; sampling from the resultant MPS would take ~ 0.3 core-seconds per sample (scaling as $\mathcal{O}(N\chi^2)$). For comparison, the experimental cycle time is ~ 1.7 s, limited by array loading and imaging; the actual quantum simulation time is only $\sim 1 \mu\text{s}$ per shot. While the classical computer can utilize multiple cores, so too can the experiment be parallelized over multiple atom-array chains simultaneously, which in fact we do already at small system sizes.

We predict these classical costs are highly-sensitive to the effective per-atom fidelity, \mathcal{F} , defined by $\mathcal{F}^{Nt} \equiv F(N, t)$ (Fig. 10.1b). For instance, the simulation time scales as $\sim (1 - \mathcal{F})^{-10}$ around the experimental \mathcal{F} . While specialized classical hardware [389, 427] may more readily perform the present approximate classical simulations, we thus expect small improvements in the quantum fidelity may soon make the experiment out of reach of even these more advanced classical systems.

For instance, using erasure conversion (Chapter 3) techniques we would expect roughly a 50% increase in fidelity at $N = 60$ from both eliminating preparation and bright state decay errors. Further, by post-selecting benchmarking purely in the blockaded Hilbert space (Chapter 8.3, roughly a further 50% gain in fidelity could be achieved, meaning the fidelity would be increased to ≈ 0.4 , or an effective per-atom fidelity of $\mathcal{F} \approx 0.999$. This would entail (Fig. 10.1 a χ^* of a few tens of thousands, simulation times on the order of thousands of core-days, sampling times exceeding the experimental duty cycle (in core-time), and terabytes of required classical RAM memory.

10.2 SIMULATION METHODS

Here we describe the various simulation methods we compared against our quantum simulator, including tensor network, path integral, neural net, and effective Hamiltonian models.

10.2.1 Description of the Caltech computing cluster

This work extensively utilizes the Caltech Resnick High Performance Computing (HPC) Center, running CentOS Linux 7. In our work, each simulation is assigned to a single node, using 16 cores of Intel Cascade Lake CPU (at 2.2 or 2.4 GHz), with a total node memory of 386 GB (at 2933 MT/s). We note that many classical simulation times are reported as core-time, which we simplify as the wall time multiplied by the number of cores.

In terms of raw wall time, the largest simulation we performed (for $N = 60$ and $\chi = 3072$) took ~ 8.3 days (~ 11.1 days if including the additional time points past the last experiment

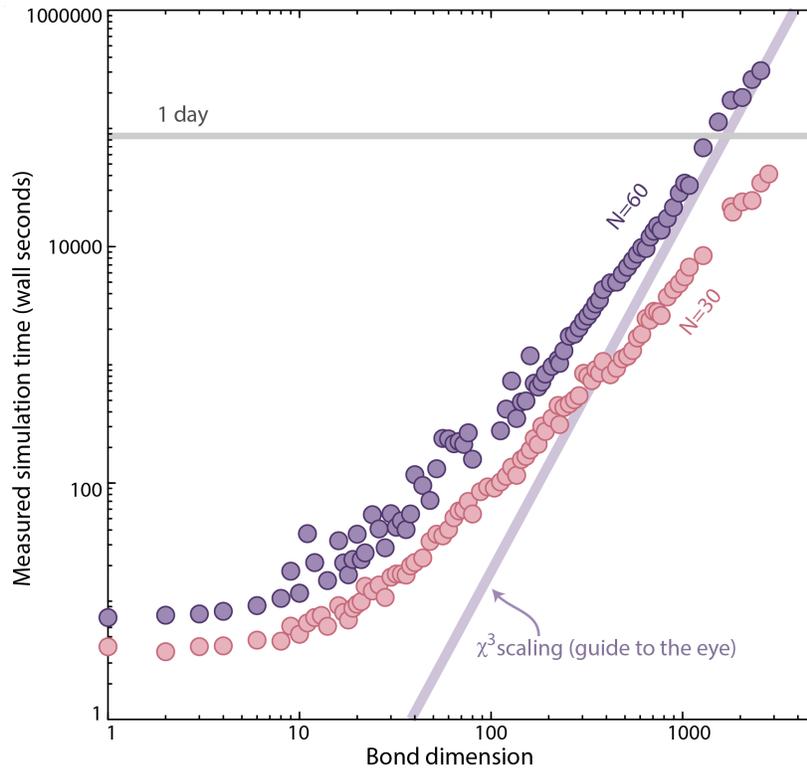


Figure 10.2: **Simulation time as a function of bond dimension.** After an initial overhead plateau, the Lightcone-MPS simulation time (measured in raw wall time on the Caltech central cluster) to reach the last experimental time increases as a function of N , approaching the expected asymptotic χ^3 scaling.

time in order to calculate p_{avg}). Representative times to simulate up to the last experimental time point are shown as a function of bond dimension in Fig. 10.2; extrapolating the expected $\mathcal{O}(\chi^3)$ scaling we predict simulation for $N = 60$ and $\chi^* = 3400$ would entail a simulation time of 11.3 wall-days (180 core-days). In total, we generated roughly ~ 100 TB of data in order to create and store the MPS tensors at various bond dimensions used in this work.

10.2.2 Schrieffer-Wolff transformations

We first consider the possibility of exact brute-force simulation of our system. The current state-of-the-art in exact RUC simulation is 45 qubits [428], while the state-of-the-art for continuous Hamiltonian time evolution is 38 qubits [389]. While clearly neither of these approaches would work natively for the largest experimental system sizes of $N = 60$, it is possible they could become applicable through a Schrieffer-Wolff transformation by rewriting the Rydberg Hamiltonian only in the blockaded subspace, applying appropriate

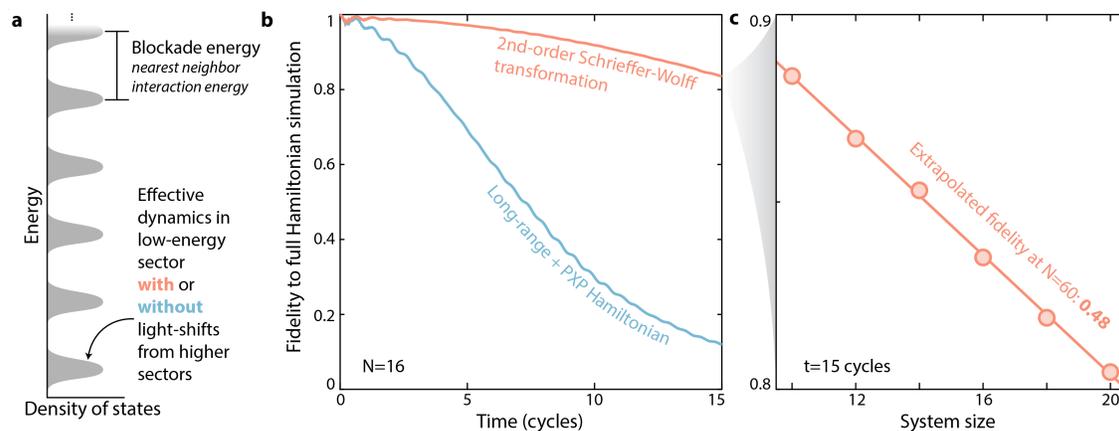


Figure 10.3: Using an effective Hamiltonian simulation. **a.** The Rydberg blockade constraint stratifies the energy levels of the Rydberg Hamiltonian, enabling approximating the dynamics in the full Hilbert space by an effective set of dynamics in the lowest energy Hilbert space (relative size of energy sectors are not to scale). **b.** While a lowest-order approximation (blue) quickly causes the fidelity of the reduced-Hilbert-space simulation to drop compared to simulation in the full space, adding the second-order term from a Schrieffer-Wolff Hamiltonian transformation [97] (orange) improves the fidelity. **c.** Even with the second-order correction, the simulation fidelity decreases as a function of system size. It is possible that numerically solving for higher order terms in the Schrieffer-Wolff transformation will improve this fidelity, but it likely can never achieve unity given the finite blockade violation probability (Fig. 8.8d).

level shifts to mimic the effect of higher blockade sectors, as described in the supplement of Ref. [97]. Because of the Rydberg blockade constraint, the effective system size is given by $\tilde{N} = \log_2(\text{Fib}(N + 2)) \approx 0.7N$, where Fib is the Fibonacci function. For the 60-atom system, $\tilde{N} \sim 42$ qubits, so such brute-force techniques could allow for potentially pseudo-exact simulation.

However, here we discount such approaches for multiple reasons. First they both rely on specialized hardware – a supercomputing system [428] with 100s of TB of RAM, or tensor processing units [389] which are not widely available. Further, even with this hardware it is not clear that our system would be easily simulable, because of difficulties associated with treating the blocked subspace. In particular, applying local gate operations becomes more complicated to implement efficiently, because the blocked subspace is not factorizable into the product of local bases.

Finally, applying this transformation is still an approximation to the true dynamics, for which it is not clear if the resulting fidelities will actually be competitive at these large system sizes and late evolution times. In Fig. 10.3 we test the accuracy of the Schrieffer-Wolff transformation at second order for system sizes up to $N = 20$ ($\tilde{N} \sim 14$), where

we see the fidelity falls by the latest experimental time. Importantly, the fidelity decay increases with system size. A simple linear extrapolation implies a simulation accuracy of ~ 0.48 for $N = 60$. Further, employing this approach will always have some fundamental inaccuracy as the blockade constraint is imperfect, and we expect the ideal target state to always have some finite population in the blockade-violating sectors, a symptom which grows stronger with increasing system size and evolution time (Fig. 8.8d). However, we do not discount the possibility that by adding successively higher-order terms in the Schrieffer-Wolff transformation, the accuracy of this method could be at least improved.

We strongly emphasize that storing even a single copy of the wavefunction at $\tilde{N} \sim 42$ still requires ~ 65 TB of memory, and that $\tilde{N} \sim 42$ exceeds the current state-of-the-art for Hamiltonian simulation [389], $N = 38$. Still, we welcome collaboration on studying the boundaries of nearly-exact classical simulation with the most advanced hardware.

10.2.3 Standard Trotterization-based TEBD-MPS

A standard approach for simulating large one-dimensional (1D) quantum systems is to use a matrix product state (MPS) representation [393]. In this representation, the quantum state is decomposed into a tensor network in the form of a 1D chain through repeated Schmidt decompositions. The key insight of such schemes is that if particles are weakly entangled at any bipartition of the system, then the on-site tensors will be low rank, and so can be efficiently *truncated* to a smaller size while preserving most of the information. The size to which the tensor is truncated is user-controllable, and is known as the *bond dimension*, χ . Thus, larger bond dimensions equate to truncating less information when converting a state into its MPS representation. For a more general review of MPS, see Ref. [392].

An important aspect of the tensor decomposition of the MPS is that unitary evolution can be efficiently applied to just part of the tensor network. This enables the standard time-evolving block decimation (TEBD) algorithm [423], in which Hamiltonian evolution is discretized (or *Trotterized*) using the Suzuki–Trotter formula

$$e^{-i(\sum_i \hat{H}_i)\delta} = \prod_i e^{-i\hat{H}_i\delta} + \mathcal{O}(\delta^2). \quad (10.1)$$

For a many-body Hamiltonian with (geometrically) local interactions \hat{H}_i , each term is applied sequentially onto the MPS. With suitable implementation, Trotterization results in a second-order decomposition of the time evolution operator.

Thus, the basic algorithmic loop of the TEBD approach is: 1) perform Schmidt decompositions to write the state as a tensor network, 2) *truncate* the tensor ranks, 3) *Trotterize* the

Hamiltonian, 4) and apply it to the relevant tensors, after which we repeat step 1. The two fundamental sources of error in this approach come from steps 2 (the truncation) and 3 (the Trotterization) (Fig. 10.5).

The error of the Trotterization decomposition depends on the product of the commutator of the terms in the Hamiltonian and the time step dt (here dt is defined as a full cycle of application of local evolution from the leftmost site to the rightmost site and then back to the leftmost site). For the Rydberg system, the commutator is on the order of the nearest neighbor interaction V_{nn} . Therefore, even with the second-order Suzuki–Trotter decomposition, the perturbative expansion converges only for $V_{nn} \cdot dt \lesssim 1$. Due to the Rydberg blockade, the nearest neighbor interaction for Rydberg Hamiltonian is strong, putting a significant constraint on the maximum dt possible for the simulation. In practice, with a second-order decomposition of the time-evolution operator, we found that a maximum of $dt \approx 1/100$ Rabi cycle is necessary to have a 0.5% Trotterization error.

Importantly, however, for too small of a timestep, the TEBD algorithm will perform many more truncations, each time lowering the fidelity. This leads to a fundamental tension, where an optimal dt must be chosen which balances these competing effects of truncation and Trotterization errors (Fig. 10.5d).

Moreover, a truncation of the long-range interaction must be chosen to efficiently use the TEBD algorithm. In practice, we found that keeping an interaction length of 5 results in the desired error of less than 0.5% (Fig. 10.4c). This results in a local evolution involving at most 6 atoms.

We emphasize that we do not believe this performance could be strongly improved through the use of a Schrieffer-Wolff transformation as described above. This is because in performing the Schmidt decompositions, MPS-based methods already find the optimal working basis, which is nearly entirely within the blockaded Hilbert space (Fig. 8.8d). In fact, it is likely such a transformation will *decrease* the MPS performance, because of the added complexity in local gate application due to the non-locally-factorizable structure of the blockaded subspace.

10.2.4 The Lightcone-MPS algorithm

One of the challenges of the Trotterization TEBD algorithm is the small dt required to faithfully simulate the quantum dynamics. However, this difficulty can be circumvented by efficiently utilizing the lightcone dynamics of quantum systems, allowing for simulation with much larger dt , and thus higher fidelity as truncation errors are minimized. In particular,

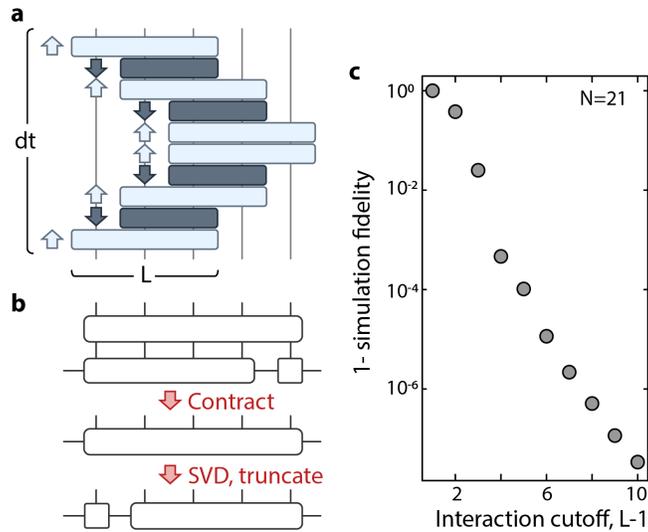


Figure 10.4: **Overview of the Lightcone-MPS algorithm.** **a.** Illustration of the Lightcone-MPS Hamiltonian decomposition over a single time step dt . The long range Rydberg interaction is truncated after a cutoff distance of L . **b.** Illustration of the MPS update procedure. **c.** We choose the Lightcone-MPS block size to be $L = 6$, which we estimate as introducing an error of $< 0.5\%$ at $N = 60$.

in Ref. [396] it was shown that a global time evolution operator could be decomposed into small blocks of local evolution by combining forward and backward time evolution. In this proposed non-perturbative decomposition, the time step is only limited by the Lieb-Robinson velocity v and the block size L of the local time evolution, not the blockade interaction strength [429].

Here, we propose a modified version of the Hamiltonian decomposition in Ref. [396], which is especially amenable to efficient integration with MPS state representations, and which we call Lightcone-MPS. This modification allows for a smaller block size L with the same cutoff of long-range interactions. This algorithm consists of alternating forward-time and backward-time evolution, where the forward-time evolution has a block size L , and the backward evolution has a block size of $L - 1$. Each block performs the time evolution for all Hamiltonian terms strictly inside the block. Crucially, in this implementation, dt is no longer limited by blockade interactions at the edges of the local evolution blocks (Fig. 10.5d).

Specifically, for the Rydberg Hamiltonian on N atoms

$$\hat{H}/h = \Omega \sum_{i=1}^N \hat{S}_i^x - \Delta \sum_{i=1}^N \hat{n}_i + \frac{C_6}{a^6} \sum_{i=1}^N \sum_{j=i+1}^N \frac{\hat{n}_i \hat{n}_j}{|i-j|^6}, \quad (10.2)$$

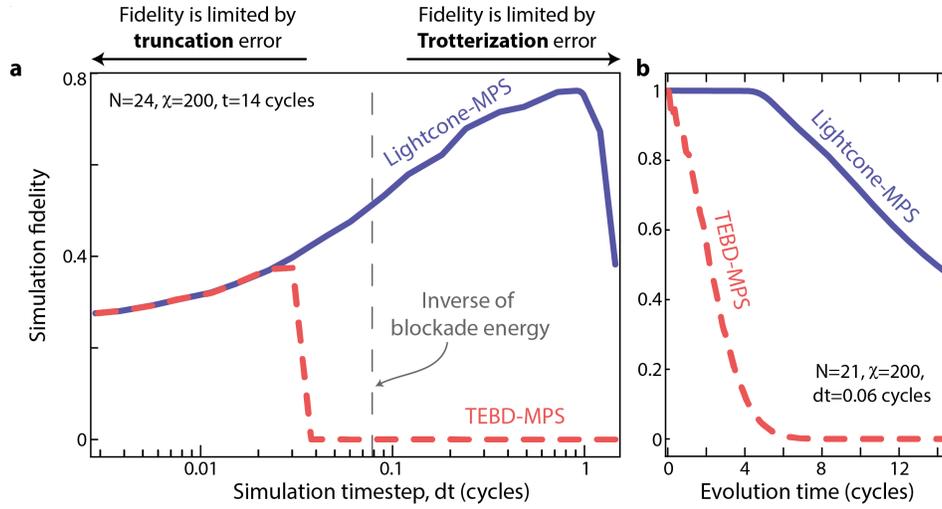


Figure 10.5: **Comparing Lightcone-MPS and TEBD algorithms.** **a, b.** For MPS algorithms there exists an optimal simulation timestep; for too large a timestep, the Hamiltonian Trotterization introduces simulation errors. For too small a timestep, the MPS is contracted and truncated too often, leading to loss of accuracy. For the TEBD algorithm, the maximum timestep is set by the inverse blockade energy, but the Lightcone-MPS algorithm can be used accurately with much larger timesteps. This ultimately means that at equal bond dimension, χ , the Lightcone-MPS algorithm can be optimized to higher fidelities than the TEBD.

we consider the Hamiltonian defined locally from k to $k + L - 1$ sites as

$$\hat{H}_{k,k+L-1}/h = \Omega \sum_{i=k}^{k+L-1} \hat{S}_i^x - \Delta \sum_{i=k}^{k+L-1} \hat{n}_i + \frac{C_6}{a^6} \sum_{i=k}^{k+L-1} \sum_{j=i+1}^{k+L-1} \frac{\hat{n}_i \hat{n}_j}{|i-j|^6}, \quad (10.3)$$

with the forward evolution operator on L sites and backward evolution operator on $L - 1$ sites defined as

$$\hat{U}_{k,k+L-1}^{\text{forward}} = \exp\left(-i \frac{\hat{H}_{k,k+L-1}}{h} \frac{dt}{2}\right); \quad (10.4)$$

$$\hat{U}_{k,k+L-2}^{\text{backward}} = \exp\left(+i \frac{\hat{H}_{k,k+L-2}}{h} \frac{dt}{2}\right) \quad (10.5)$$

Then, the evolution within one timestep is decomposed (Fig. 10.4a) as

$$\exp\left(-i \frac{\hat{H}}{h} dt\right) \approx \hat{U}_{1,L}^{\text{forward}} \prod_{i=2}^{N-L+1} \left(\hat{U}_{i,i+L-2}^{\text{backward}} \hat{U}_{i,i+L-1}^{\text{forward}}\right) \prod_{i=N-L+1}^2 \left(\hat{U}_{i,i+L-1}^{\text{forward}} \hat{U}_{i,i+L-2}^{\text{backward}}\right) \hat{U}_{1,L}^{\text{forward}}. \quad (10.6)$$

Here, the order of applying the time evolution operator is reversed in every other layer (i.e. left to right for odd layers and right to left for even layers), and each two layers are

defined as one timestep. In practice, each forward evolution can be combined with the backward evolution (each pair of operators in Eq. (10.6)) into a single operator, reducing the simulation cost.

The simulation cost can be further reduced by always keeping a tensor with $L - 1$ atoms fused together, without performing addition decompositions of the MPS. (Fig. 10.4b). More specifically, we can keep an MPS in the form

$$\langle z_1, z_2, \dots, z_N | \psi_{MPS} \rangle = \sum_{\alpha_1, \alpha_2, \dots, \alpha_{N-1}} A_{z_1}^{\alpha_1} A_{z_2}^{\alpha_1 \alpha_2} \dots M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}} \dots B_{z_N}^{\alpha_{N-1}}, \quad (10.7)$$

without decomposing $M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}}$ into $L - 1$ tensors. Here, z_i denotes the physical indices and α_i denotes the bond dimensions. In addition, we use A to denote the tensors in the left canonical form, B to denote tensors in the right canonical form, and M to denote tensors in the center canonical form. Basically, all tensors to the left of $M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}}$ should be in the left canonical form and all the tensors to the right of $M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}}$ should be in the right canonical form. When applying the operator of length L between site j and site $j + L - 1$, we can first combine $M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}}$ with $B_{z_{j+L-1}}^{\alpha_{j+L-2} \alpha_{j+L-1}}$ into a tensor of length L

$$M_{z_j, \dots, z_{j+L-1}}^{\alpha_{j-1} \alpha_{j+L-1}} = \sum_{\alpha_{j+L-2}} M_{z_j, \dots, z_{j+L-2}}^{\alpha_{j-1} \alpha_{j+L-2}} B_{z_{j+L-1}}^{\alpha_{j+L-2} \alpha_{j+L-1}}. \quad (10.8)$$

Then, we update the tensor by applying the operator of length L

$$M'_{z_j, \dots, z_{j+L-1}}^{\alpha_{j-1} \alpha_{j+L-1}} = \sum_{z'_j, \dots, z'_{j+L-1}} U_{z_j, \dots, z_{j+L-1}; z'_j, \dots, z'_{j+L-1}} M_{z'_j, \dots, z'_{j+L-1}}^{\alpha_{j-1} \alpha_{j+L-1}}, \quad (10.9)$$

and decompose the updated tensor back into two tensors of length 1 and $L - 1$

$$M'_{z_j, \dots, z_{j+L-1}}^{\alpha_{j-1} \alpha_{j+L-1}} \xrightarrow{SVD} A'_{z_j}^{\alpha_{j-1} \alpha_j} M'_{z_{j+1}, \dots, z_{j+L-1}}^{\alpha_j \alpha_{j+L-1}}. \quad (10.10)$$

Here, the process of truncating the Schmidt coefficient and combining it into $M'_{z_{j+1}, \dots, z_{j+L-1}}^{\alpha_j \alpha_{j+L-1}}$ is omitted for simplicity. The same process can be carried on over the rest of the system, and reversed if the operators are applied in the reverse order.

The Lightcone-MPS algorithm allows for a maximum $dt \sim L/v$. Although computing the Lieb-Robinson velocity is difficult, based on empirical results, we find in practice that $dt \approx 0.5$ Rabi cycle is sufficient for a 0.5% decomposition error with a block size of 6 (long-range interaction cutoff at 5). Again, the dt is defined as a full cycle of application of local evolution from the left-most site to the right-most site and then back to the left-most site.

Crucially, there is still a tension on dt from the competing effects of truncation and Trotterization errors, but for the Lightcone-MPS algorithm the Trotterization errors are minimized up to much larger dt . This means at equal bond dimension, if selecting the optimal dt , the Lightcone-MPS can reach higher fidelities than TEBD (Fig. 10.5de). Equivalently, to reach the same fidelity necessitates the TEBD to use a larger bond dimension, as is visible in Fig. 10.1.

Additionally, even if a larger bond dimension is used for the TEBD (where truncation error is small), the optimal dt in TEBD is empirically $\sim 50\times$ smaller than that for the Lightcone-MPS, increasing the raw simulation time further. In practice, to benchmark the experiment, we need to choose dt such that we can match the time step chosen in the experiment (around 0.77 Rabi cycle). Therefore, we use half the experimental time step as the dt for the Lightcone-MPS algorithm, and use $1/80$ the experimental time step as the dt for the TEBD algorithm. This still results in a $1/40$ reduction in number of simulation steps thanks to the Lightcone-MPS algorithm.

10.2.5 Time-dependent variational principle

The Time-Dependent Variational Principle [425] (TDVP) algorithm is another MPS time evolution algorithm. In this algorithm, the Schrödinger equation is projected onto the variational manifold of MPS, where an effective (nonlinear) partial differential equation (PDE) is obtained to update the MPS parameters. This method is better than TEBD in certain situations such as simulating emergent hydrodynamics, due to its ability to preserve all symmetries and conservation laws. In addition, TDVP can naturally handle long-range interactions.

However, the TEBD algorithm (and the Lightcone-MPS algorithm) is expected to be better suited in our setup. First, the TEBD algorithm is optimal for minimizing the infidelity in each evolution time step. This is done by truncating the Schmidt coefficients in a way that the infidelity is minimized. In the TDVP algorithm, on the other hand, the infidelity is not necessarily minimized at the cost of preserving conserved quantities. Second, when combined with MPS representation, the TDVP implementation is nontrivial. This is due to the existence of the gauge degree of freedom in MPS representation. More specifically, the representation of MPS is not unique due to gauge transformation. This makes it difficult to utilize MPS-based TDVP by leveraging existing PDE solvers. Ref. [430] presents a solution to this problem, which shows that TDVP can be implemented by using the canonical form of MPS, but it relies on evolving the system locally in small time steps, akin to the Trotterization of the Hamiltonian. In addition, it turns out that this solution is a minor modification of

the original TEBD algorithm, where a backward time evolution is added. This backward evolution is already naturally incorporated in our Lightcone-MPS algorithm.

Moreover, regardless of the implementation of the TDVP algorithm, the PDE needs to be solved numerically, with the maximum time step limited by the largest term of the Hamiltonian or the largest term in the PDE. Therefore, the TDVP algorithm has the same bottleneck as the TEBD algorithm: the time step must be small compared to the blockade interaction strengths, limiting its utility. In our Lightcone-MPS algorithm, however, this limitation is lifted by utilizing extra physical properties such as Lieb-Robinson bound and our knowledge about the geometry of the system (Fig. 10.4). Given this, we do not explicitly test a TDVP implementation in this work.

10.2.6 Noisy MPO-based simulation

Instead of simulating perfect quantum dynamics with limited fidelity, an alternative method of benchmarking with the experiment is to (almost) exactly simulate noisy quantum dynamics $\hat{\rho}(t)$ with similar fidelity. We explore this possibility using the recently developed matrix product operator (MPO) based simulation [424]. In MPO simulation, each local tensor has an extended Hilbert space dimension of 4 instead of 2. The increased local dimension significantly increases the cost of simulation using the efficient decomposition of Ref. [396]. Therefore, we fall back to simulating the Trotterized Hamiltonian dynamics and ignore the long-range interactions. In addition, we choose dt to be half the Rabi cycle, and ignore Trotterization error, in order to consider this method in the most favorable light.

For simplicity, we consider local dephasing error as the only error source and find the effective dephasing rate γ per qubit as a function of system size such that the fidelity of the simulation matches the fidelity of the experiment (Fig. 10.6), extrapolated using the exponential decay fit. As shown in the figure, the effective dephasing rate decreases as the system size increases, which means the dephasing rate at $N = 12$ overestimates the dephasing rates at larger system sizes. Since the MPO entanglement decreases as the decoherence rate increases, this leads to an underestimation of the MPO simulation complexity. To make the case more favorable for the MPO method, we thus use the dephasing rate at $N = 12$ for the simulation at $N = 60$.

Assuming the early time exponential fit of the experimental fidelity, the Lightcone-MPS algorithm achieves similar fidelity with the experiment at $N = 60$ using $\chi = 913$ (Fig. 10.10). This motivates us to test if the MPO method with a similar or larger bond dimension can achieve the same fidelity. We simulate the noisy quantum dynamics at $N = 60$ using

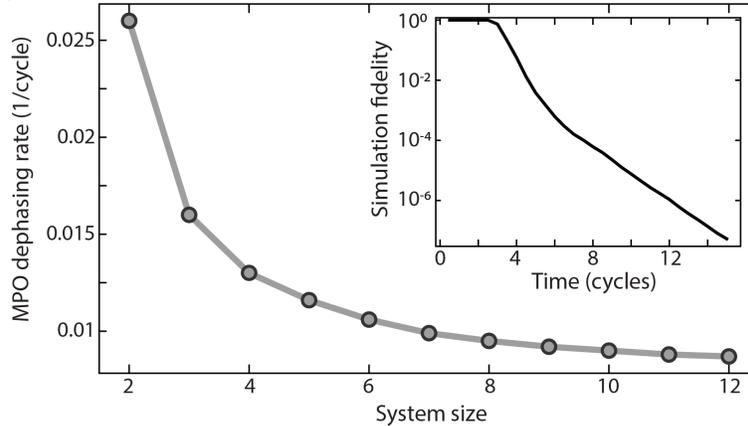


Figure 10.6: **Matrix product operator (MPO) simulation.** Effective dephasing rate per qubit to obtain noisy dynamics with a fidelity similar to the exponential decay fit of the experiment. Inset) Truncation accuracy of the MPO method with $\chi = 1100$ at $N = 60$. The effective dephasing rate is taken as the dephasing rate at $N = 12$, which should serve as an upper bound of the effective dephasing rate at $N = 60$.

$\chi = 1100$ and the dephasing rate upper bound, yielding the un-normalized simulation mixed state $\hat{\rho}_{trun}$, where we do not explicitly normalize $\hat{\rho}_{trun}$ at each time step. Here, $\text{tr}(\hat{\rho}_{trun})$ serves as a measure of the simulation accuracy, where it equals 1 when the simulation is perfect and decreases as we truncate the density matrix. We measure $\text{tr}(\hat{\rho}_{trun})$ as a function of time (Fig. 10.6, inset) and find that it drops to $\sim 10^{-7}$. This shows that the MPO algorithm is far from an exact simulation of the noisy quantum dynamics. We note that the $\text{tr}(\hat{\rho}_{trun})$ measures the accuracy of the algorithm compared to the noisy dynamics. It does not measure the fidelity of the simulation compared to the ideal dynamics, which can be even smaller. Thus, this study, in combination with the Trotterization errors not considered here, implies the MPO-based approach will achieve far lower fidelity than the experiment and the lightcone-MPS method for reasonable simulation parameters.

10.2.7 Path integral formulation

A hybrid path integral algorithm was introduced in Ref. [162] for quantum circuit simulation. In this algorithm, the system is cut into two patches. When no gate is applied between the two patches, each patch evolves independently under the Schrödinger equation. When a gate is applied across the two patches, the gate is decomposed into sums of gates that apply independently on the two patches, where one path was sampled out of this decomposition. More formally, we write the gate as $\hat{U} = \sum_i \hat{V}_i \otimes \hat{W}_i$. For each path, we only choose a particular $\hat{V}_i \otimes \hat{W}_i$, so that the two patches stay in a product state. The wave function of

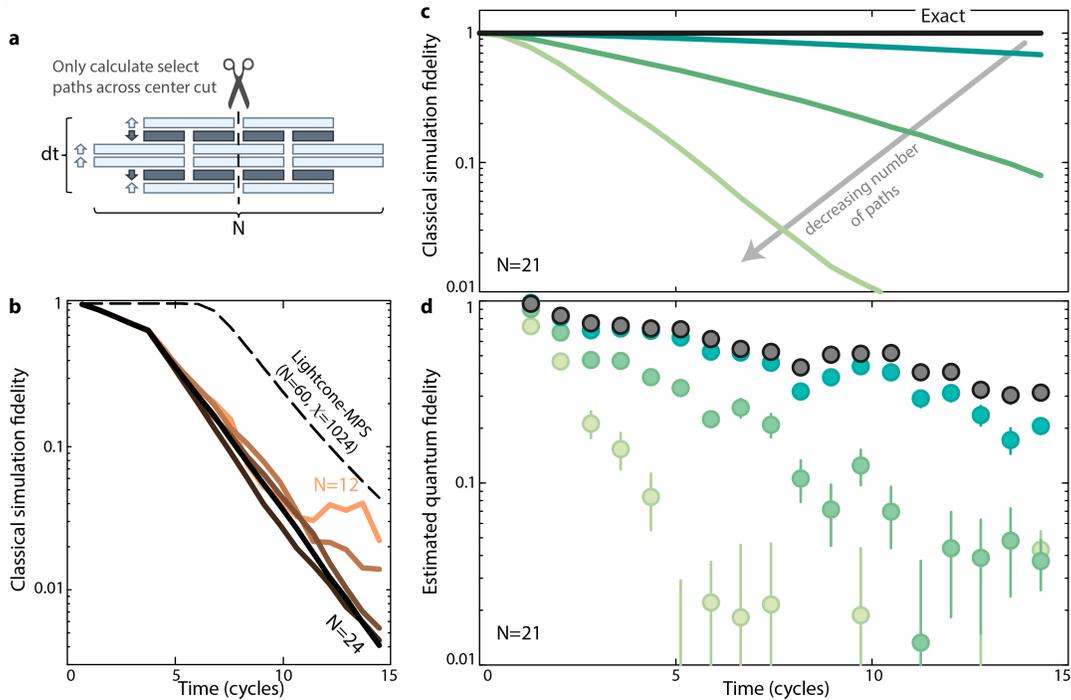


Figure 10.7: **A hybrid path integral simulation.** **a.** Illustration of the hybrid lightcone-path integral (Lightcone-PI) method. **b.** Simulation fidelity of the Lightcone-PI with 4×10^6 paths, showing decreasing fidelity as a function of N . We compare against F_{svd} for the Lightcone-MPS algorithm at $\chi = 1024$ and $N = 60$, which outperforms the L-PI method. **c, d.** At a fixed N , the infidelity in the Lightcone-PI method from using too few paths is visible when benchmarking directly against experimental data, as was the case for the Lightcone-MPS algorithm (Fig. 8.1).

each path can be then summed over to obtain the desired result. Suppose each gate across the cut can be minimally decomposed as r product gates and there are m gates across the cut in total, r^m paths are required to fully describe the system.

In this work, we adapt the hybrid path integral algorithm to the long-range Hamiltonian dynamics. We first decompose [396] the global time evolution operator into local gates with block size $L = 10$. We further reverse the order of the gate application at every other layer so the gates can be fused between the two layers (Fig. 10.7). This reduces the number of gates that need to be decomposed at the cut. This allows us to have a minimum of 19 cuts throughout the simulation to match the experiment time steps. Similar to the original hybrid path integral algorithm, we cut the system into two halves. For gates applied across the two halves of the system, we use Schmidt decomposition to write the gate as $\hat{U} = \sum_i \sigma_i \hat{V}_i \otimes \hat{W}_i$, where σ_i can be then absorbed into either of the halves. However, a naïve decomposition results in a large interaction between the blockade and non-blockade

sectors, which results in poor convergence. Therefore, we project the gate into the blockade sector before performing the decomposition.

We test this algorithm for up to 24 atoms where we truncate the rank of the gates across the cut to r_i to mimic actually simulating the path integral with $\prod_i r_i$ paths. We find that the resulting fidelity mostly depends on the number of paths and only weakly depends on the system size, where the dependency becomes weaker as the system size increases. This is because, for all system sizes, we have the same number of gates across the cut, which we expect to be the only factor affecting the fidelity for large enough system sizes. Extrapolating to 60 atoms, we find that we need at least 4×10^6 paths to obtain a fidelity of ~ 0.01 , where *each path* requires an exact simulation of two 30-atom systems. Even using our Lightcone-MPS method, simulating each 30-atom system takes 60 core-hours (with $\chi = 1536$, which is nearly exact). This adds up to a formidable amount of computational resources. In comparison, a direct application of the Lightcone-MPS algorithm for 60 atoms achieves better fidelity with a bond dimension of $\chi = 1024$ (Fig. 10.7b).

Interestingly, we can still benchmark the experimental fidelity using the hybrid path integral approach, which we see shows qualitatively the same behavior as the Lightcone-MPS algorithm in Fig. 8.1. That is, when the number of paths is too low, the classical simulation fidelity drops, and this drop is visible directly on the benchmarked experimental fidelity (Fig. 10.7cd).

This raises the interesting potential prospect of using the experimental benchmarking protocol in reverse to determine the fidelity of the classical algorithm, which we leave to future work.

10.2.8 Neural network based Hamiltonian simulation

Neural networks have shown great potential in simulating quantum dynamics [426, 431]. However, such protocols are still in an early age of development and their application to the Rydberg Hamiltonian with ergodic, long-time dynamics is still challenging. We considered two different neural network algorithms used in Refs [432, 433], but neither results in a fidelity higher than the Lightcone-MPS algorithm, even for comparing both at system sizes as low as $N = 12$. While it is possible that neural network based algorithms require additional fine-tuning and hyper-optimization, as it is not the main focus of this work, we leave them for future study. We also note that neural network algorithms are under active development, and we welcome readers to benchmark their algorithms on our data.

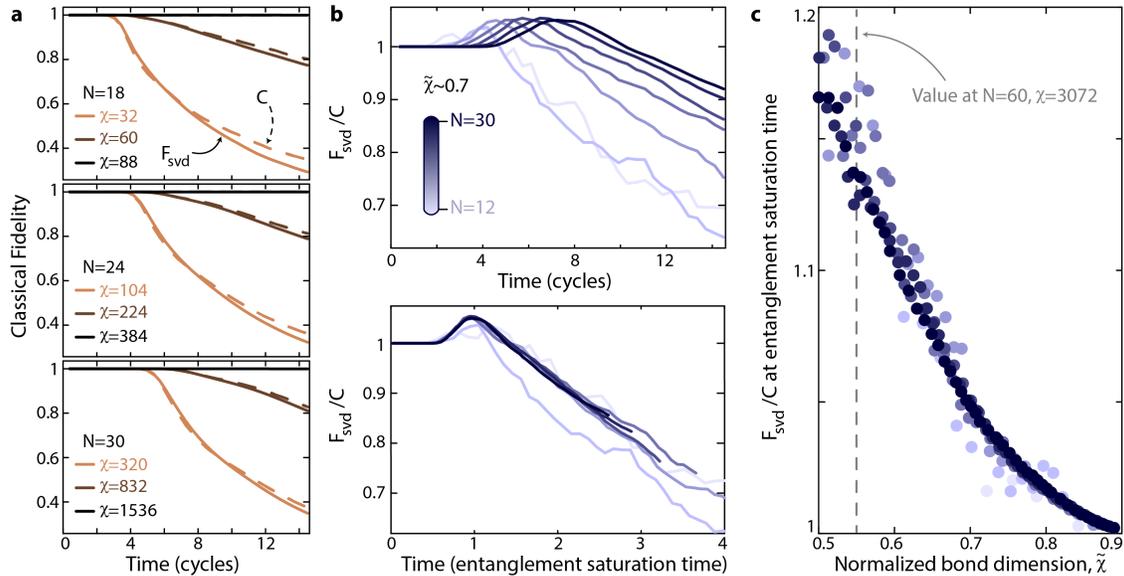


Figure 10.8: Estimating the classical fidelity. **a.** MPS truncation accuracy, F_{svd} , and MPS fidelity, C , for various bond dimensions and system sizes. The MPS fidelity is measured against another MPS calculation with $\chi = 3072$, which is almost exact for $N \leq 30$. **b.** For a fixed log-normalized bond dimension, $\tilde{\chi} \sim 0.7$, F_{svd} first over-estimates C , before under-estimating it at late time (left). However, rescaling the time by the entanglement saturation time (which is linearly proportional to system size, see Fig. 8.3e) causes the data to collapse (right). This implies that before t_{ent} , F_{svd} will over-estimate C , so in the parameter regimes we study in this work, using F_{svd} is the conservative choice when comparing quantum and classical fidelities. **c.** The ratio of F_{svd}/C around the entanglement saturation time, showing the over-estimation increases with decreasing $\tilde{\chi}$. Given the data collapse, we expect that F_{svd} overestimates by a factor of 1.15 at the entanglement saturation time for $N = 60$. Taking this effect into account implies the actual $\chi^*(N=60) \approx 4000$, higher than the value of 3400 quoted earlier in this Chapter. However, we choose to present the lower value as a more conservative estimate. We note that this relationship between F_{svd} and C is only observed to hold for the case of the Lightcone-MPS algorithm, because the timestep is long enough that truncations are relatively independent. The agreement is not observed with a TEBD MPS implementation (which requires much finer timesteps, see text).

10.3 ESTIMATING THE MPS FIDELITY

While MPS gives a controlled approximation of the exact wave function, an exact calculation of the MPS fidelity, C – defined as the overlap between the MPS state and the ideal target state – is not possible over multiple time steps for system sizes greater than $N \sim 30$. However, it is possible to estimate the MPS fidelity from the truncation fidelity

$$F_{\text{svd}} = \prod_i \sum_{\alpha=1}^{\chi} s_{i,\alpha}^2, \quad (10.11)$$

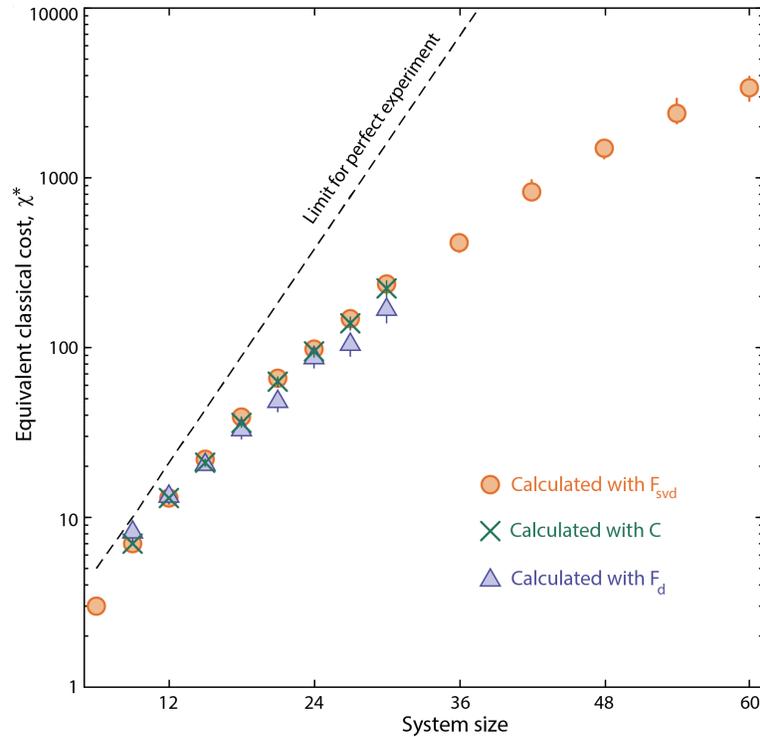


Figure 10.9: **Estimating χ^* from F_d and true classical fidelity.** For system sizes larger than $N > 30$, we are not easily able to calculate the exact classical MPS fidelity (as it is impractical to simulate the exact reference state), so we have used F_{svd} to compare against the experimental quantum fidelity. However, for $N \leq 30$, we can directly compare against the classical MPS fidelity, C , which we find gives consistent values of χ^* in the applicable range. Finally, for these small N , we can directly calculate F_d between an approximate MPS simulation and an exact simulation. We find the corresponding χ^* is also consistent with the values based solely on fidelity.

where i runs over all steps involving Schmidt value truncations, and $s_{i,\alpha}$ are the Schmidt values at truncation step i . Here, we assume the wave function is normalized where $\sum_{\alpha=1}^{\infty} s_{i,\alpha}^2 = 1$. While this estimation is only approximate, it can be extremely accurate when successive truncations are independent [394]. For $N \leq 30$, a bond dimension of $\chi = 3072$ almost exactly captures the dynamics, therefore, it is possible to compute the MPS fidelity, C , for various bond dimensions. In Fig. 10.8a, we plot F_{svd} versus the actual MPS fidelity, C , to show the general agreement when using the Lightcone-MPS algorithm. This is expected due to the large dt of the algorithm, which makes successive truncations approximately independent.

In addition, we notice that the agreement becomes better when the system size increases, indicating that F_{svd} can be a good estimator of the MPS fidelity for large systems (Fig. 10.8b). If anything, it appears that in the parameter regime of evolution times shorter than the

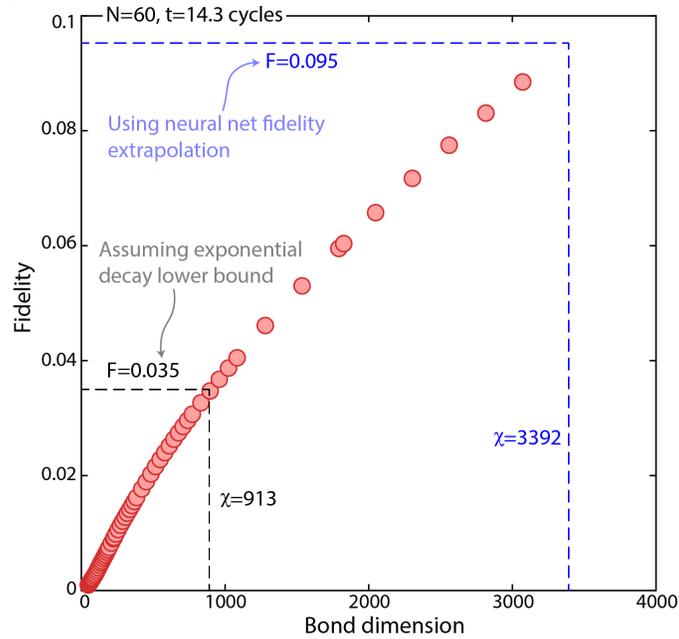


Figure 10.10: **Extrapolating χ^* .** MPS fidelity (as quantified by the MPS truncation, F_{svd}) for bond dimensions up to $\chi = 3072$, the largest directly simulated. None of the MPS fidelities (markers) exceed the experimental value (blue dashed lines), but given the functional smoothness in this regime, we use a linear extrapolation of the last few data points to predict the χ^* for the experimental fidelity, yielding $\chi^* = 3392 \approx 3400$.

entanglement saturation time, F_{svd} slightly over-estimates C , making it a conservative quantity to study (Fig. 10.8bc). Further, we emphasize that for $N < 30$, where we have access to both C and F_{svd} , the values of χ^* are consistent (Fig. 10.9).

10.3.1 Extrapolating the experimental classical cost

We only directly perform Lightcone-MPS simulations up to $\chi = 3072$. At this highest χ , the MPS fidelity for $N = 60$ at the latest experimental time is 0.088, as quantified by F_{svd} . Given that this is lower than the experimental fidelity of 0.095(11), we need to extrapolate in order to find the predicted value of χ^* . The fidelities are very close, so we make a simple linear extrapolation approximation, from which we find $\chi^* = 3392 \approx 3400$ (Fig. 10.10a).

10.3.2 Fitting and extrapolating the MPS fidelity

In the previous subsection, we used a linear extrapolation of MPS fidelity as a function of bond dimension in order to predict χ^* beyond the regime we directly simulated. However, in Fig. 10.1b, we compare the experimental and MPS fidelity over a much larger hypothetical

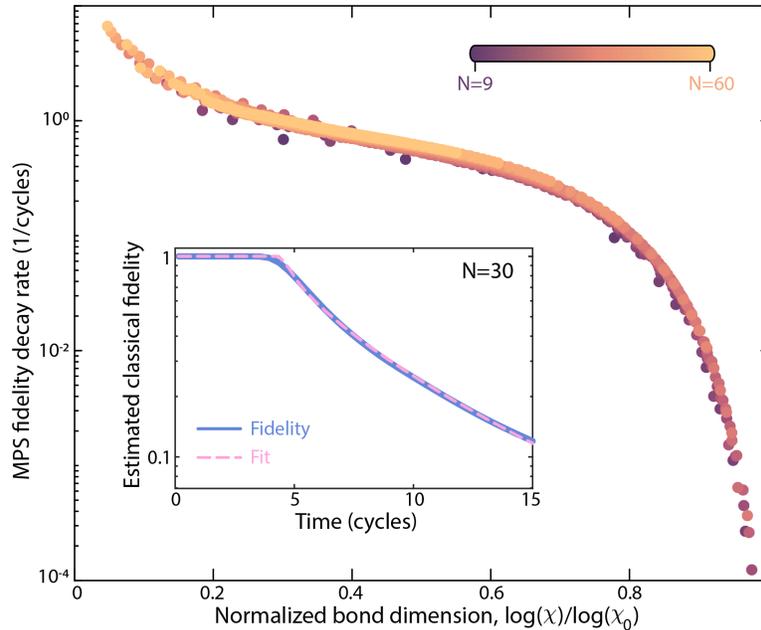


Figure 10.11: **A universal form of MPS fidelity decay.** For estimating the χ required to reach a broader range of fidelities (as for instance is used in Fig. 10.1b), we fit the Lightcone-MPS fidelity as a stretched exponential decay past the exact simulation time (inset). We observe a universal collapse of the decay rate as a function of the log-normalized bond dimension, for all system sizes and bond dimension up to $N = 60$ and $\chi = 3072$.

fidelity range, where the linear approximation breaks down (see for instance Fig. 8.16ab).

In order to make the comparisons in Fig. 10.1b, we perform fits of the MPS fidelity, approximated by F_{svd} . For RUCs, F_{svd} is essentially exponentially decaying past t_{ex} , as the truncations are independent [394]. In our case this behavior is not perfect, and so we fit the MPS fidelity as a stretched exponential decay after t_{ex} .

We are able to fit the exponential parameters, and in particular observe an apparent universality in the decay rate as a function of the log-normalized bond dimension (Fig. 10.11b), an interesting phenomena which we leave to future study. We ultimately use these fits to predict the MPS fidelity at arbitrary bond dimension to compare against experiment; we find the fits generally agree to within a few percent with the true fidelities.

To generate Fig. 10.1b, we assume an effective per-atom error rate, \mathcal{F} , and find the minimum χ^* such that the fitted MPS fidelity is always greater than $F = \mathcal{F}^{Nt}$. The experimental \mathcal{F} is the value required for exponential decay of the experiment to match the fidelity value of 0.095(11) for the experiment at $t = 14.3$ cycles and $N = 60$. We then convert χ^* into runtime using the known conversion of $\chi = 3072 \rightarrow 136$ core-days, measured empirically on the Caltech cluster, and assuming the runtime scales as $\mathcal{O}(\chi^3)$ (Fig. 10.2).

10.4 NEGATIVITY OF TRUNCATED MPS STATES

We now investigate an interesting observation, that at $N = 60$ the entanglement-proxy found for the experiment in Fig. 9.2 is ~ 3 ebits higher than $\log_2(\chi^*)$ in Fig. 10.1. In other words, the fidelity-equivalent MPS simulation is less entangled than the experiment. Here we study the fidelity and entanglement (negativity) of truncated pure state MPS representations to explain this behavior. Ultimately our results indicate there may be a deeper connection between mixed state entanglement and computational complexity, a thorough investigation of which we leave to future work.

For a theoretical treatment, we consider a simplified setup where we start with an ideal pure state and its equal bipartite Schmidt decomposition $|\psi\rangle = \sum_{j=1}^{D_A} s_j |a_j\rangle |b_j\rangle$ where D_A is the Hilbert space dimension of the half system and the Schmidt values s_j are assumed to be in descending order. Then the Schmidt rank is truncated to a maximum bond dimension $\chi \leq D_A$. Denote the truncated state and its Schmidt decomposition as $|\psi_\chi\rangle = \sum_{j=1}^{\chi} s'_j |a_j\rangle |b_j\rangle$, with $s'_j = s_j / \sqrt{\sum_{k=1}^{\chi} s_k^2}$ (normalization condition). It is straightforward to see that

$$F = |\langle \psi_\chi | \psi \rangle|^2 = \sum_{j=1}^{\chi} s_j^2. \quad (10.12)$$

To derive the entanglement negativity, the eigenvalues λ of the partial transposed state $|\psi\rangle\langle\psi|^{T_A}$ is given by s_j^2 or $\pm s_i s_j$ as mentioned above. This gives

$$\| |\psi_\chi\rangle\langle\psi_\chi|^{T_A} \|_1 = \frac{\left(\sum_{j=1}^{\chi} s_j\right)^2}{\sum_{j=1}^{\chi} s_j^2}, \quad (10.13)$$

and therefore

$$\mathcal{E}(|\psi_\chi\rangle) = \log_2 \| |\psi_\chi\rangle\langle\psi_\chi|^{T_A} \|_1 = 2 \log_2 \sum_{j=1}^{\chi} s_j - \log_2 \sum_{j=1}^{\chi} s_j^2. \quad (10.14)$$

For generic states with unknown s_j , it is hard to draw any meaningful relationship between fidelity and entanglement negativity. Therefore, we consider two cases: maximally entangled states and Haar random states.

Maximally entangled states—For a maximally entangled state, the Schmidt values are given by $s_j = 1/\sqrt{D_A}$. It follows directly that

$$F = \frac{\chi}{D_A}; \quad (10.15)$$

$$\mathcal{E}(|\psi_\chi\rangle) = \log_2 \chi = \log_2 D_A + \log_2 F, \quad (10.16)$$

with $\log_2 D_A$ the log negativity of the ideal maximally entangled pure state. Thus, the log negativity for a fidelity-equivalent truncated maximally entangled state is equal to that of an isotropic mixed state.

Haar random states—The case for Haar random states is more interesting, as such states arise from random unitary circuit evolution, and are closer to the output of our Rydberg experiment as well. As mentioned previously, the Schmidt values of Haar random states follow the *quarter-circle law* as

$$P(\tilde{s}) = \frac{1}{\pi} \sqrt{4 - \tilde{s}^2}, \quad (10.17)$$

where $\tilde{s} \equiv \sqrt{D_A} s$. Setting a maximum Schmidt rank can be translated to setting a minimum cutoff s_{\min} (and the corresponding $0 \leq \tilde{s}_{\min} \leq 2$) of the Schmidt values (keeping only Schmidt values above s_{\min}), where the relation between χ and \tilde{s}_{\min} is given by

$$\chi = D_A \int_{\tilde{s}_{\min}}^2 P(\tilde{s}) d\tilde{s} = D_A \left[\frac{4}{\pi} \sin^{-1} \left(\frac{\sqrt{2 - \tilde{s}_{\min}}}{2} \right) - \frac{1}{2\pi} \tilde{s}_{\min} \sqrt{4 - \tilde{s}_{\min}^2} \right]. \quad (10.18)$$

Depending on the cutoff,

$$\sum_{j=1}^{\chi(\tilde{s}_{\min})} s_j = \sqrt{D_A} \int_{\tilde{s}_{\min}}^2 \tilde{s} P(\tilde{s}) d\tilde{s} = \frac{\sqrt{D_A}}{3\pi} (4 - \tilde{s}_{\min}^2)^{\frac{3}{2}}; \quad (10.19)$$

$$\sum_{j=1}^{\chi(\tilde{s}_{\min})} s_j^2 = \int_{\tilde{s}_{\min}}^2 \tilde{s}^2 P(\tilde{s}) d\tilde{s} = \frac{4}{\pi} \sin^{-1} \left(\frac{\sqrt{2 - \tilde{s}_{\min}}}{2} \right) + \frac{1}{4\pi} \tilde{s}_{\min} (2 - \tilde{s}_{\min}^2) \sqrt{4 - \tilde{s}_{\min}^2}. \quad (10.20)$$

Therefore, the fidelity and log negativity can be written in parametric forms as

$$F(\tilde{s}_{\min}) = \frac{4}{\pi} \sin^{-1} \left(\frac{\sqrt{2 - \tilde{s}_{\min}}}{2} \right) + \frac{1}{4\pi} \tilde{s}_{\min} (2 - \tilde{s}_{\min}^2) \sqrt{4 - \tilde{s}_{\min}^2}; \quad (10.21)$$

$$\mathcal{E}(|\psi_{\chi(\tilde{s}_{\min})}\rangle) = \log_2 D_A + \log_2 \frac{(4 - \tilde{s}_{\min}^2)^3}{9\pi^2} - \log_2 F(\tilde{s}_{\min}). \quad (10.22)$$

It is easy to verify that when $\tilde{s}_{\min} = 0$, the equations reduce to the case of ideal Haar random states, with $\chi = D_A$, $F = 1$, and $\mathcal{E} = \log_2 D_A + \log_2 64/9\pi^2$, where $64/9\pi^2 \approx 0.72$ is the Page correction [324].

Importantly, we can then compare this prediction against the prediction for a depolarized Haar random state, where we use the first-order approximation from Eq. (9.13), in other words taking $D_A \rightarrow \infty$. At a given fidelity for each, the difference in estimated entanglement, δ , is given as

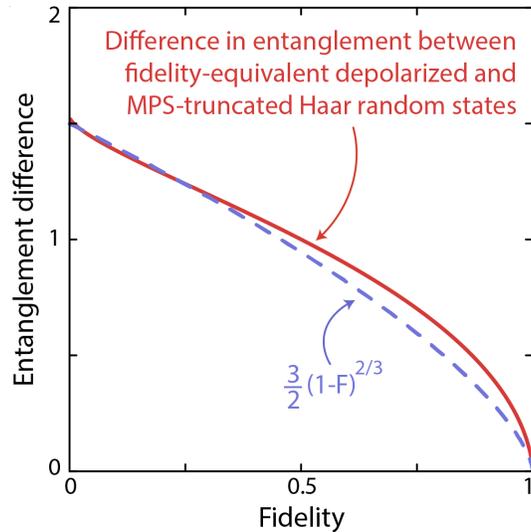


Figure 10.12: **Entanglement difference between truncated and depolarized Haar random states.** At fixed fidelity, MPS truncation and depolarization affect the entanglement content of a pure state differently. For a Haar random state, we can solve this relation analytically, finding the depolarized state has a higher entanglement (log negativity) content than the MPS-truncated state at equal fidelity.

$$\delta = \mathcal{E}_N(\hat{\rho}_{\text{Haar}}) - \mathcal{E}_N(|\psi_{\chi(\tilde{s}_{\min})}\rangle) \quad (10.23)$$

$$= 3 \log_2 \left(\frac{4}{4 - \tilde{s}_{\min}^2} \right) + 2 \log_2(F) \quad (10.24)$$

$$\approx \left(2 + \log_2 \left(\frac{64}{9\pi^2} \right) \right) (1 - F)^{2/3} \approx \frac{3}{2} (1 - F)^{2/3} \quad (10.25)$$

In the last line we have performed series expansions of both F and δ in terms of \tilde{s}_{\min} to arrive at a non-parametric form. In Fig. 10.12 we plot δ (as well as the non-parametric approximation) as a function of F , which is always positive, indicating that for all F , the MPS representation of a truncated Haar random state is less entangled (by $\mathcal{O}(1)$ ebits) than a equivalent-fidelity depolarized Haar random mixed state. Given that the experiment does not produce a perfect Haar random state, as well as the fact that a real MPS simulation performs many truncations between potentially correlated timesteps, these results are in good agreement with the experimentally found discrepancy of ~ 3 ebits. We thus conjecture there is a deeper connection between the mixed state entanglement proxy we have introduced and the classical computational complexity, though leave further investigations to future work.

Coda

As quantum systems tackle tasks of rising complexity, it is increasingly important to understand their ability to produce states in the highly entangled, beyond-classically-exact regime. Here we have studied this regime directly by measuring the global fidelity of an analog quantum simulator with up to 60 atoms, the first such demonstration to our knowledge.

A careful analysis (Fig. 8.22) indicates that with reasonable classical resources, our Monte Carlo inference protocol is scalable to an order-of-magnitude larger system sizes than were studied here, potentially enabling fidelity estimation for system sizes with $N \sim 500$. It is also applicable for digital devices [162, 316, 381, 382, 395] that are affected by non-Markovian noises like control errors [162], which could then lead to non-exponential scaling of global fidelities in certain parameter regimes. Additionally, it could be applied to analog quantum simulators for itinerant particles [30, 59, 65, 434]. Further, one may imagine applying the same basic technique to cross-platform comparisons [321, 435, 436] between erroneous quantum devices by varying the decoherence of each, a form of zero-noise extrapolation [437, 438].

In addition, we have addressed a longstanding problem by introducing a simple proxy of the experimental mixed state entanglement. This entanglement-proxy can serve as a universal quality-factor comparable amongst analog and digital quantum devices as a guide for improving future systems, and may act as a probe for detecting topological order [439, 440] and measurement-induced criticality [441].

Finally, we have studied the equivalent classical cost of our experiment on the level of global fidelity, which we note could be greatly increased through the use of erasure conversion [26, 68, 75]. Similar techniques could be applied to quantify the classical cost of measuring physical observables [66, 385], and to benchmark the performance of approximate classical algorithms themselves through comparison to high fidelity quantum data.

While here we have focused on one-dimensional systems to exploit the power of MPS representations, using higher-dimensional systems [42, 43], while maintaining high fidelities, may prove even more difficult for classical algorithms. We emphasize that in contrast to many previous experiments [162, 316, 381, 382] which explicitly targeted spatiotemporally complex quantum evolution when exploring the limits of classical simulation, here the dynamics we have studied are one-dimensional and both space- and time-independent, yet still begin to reach a regime of classical intractability.

Ultimately, our results showcase the present and potential computational power of analog quantum simulators, encouraging an auspicious future for these platforms [59].

Outlook

Over the course of this thesis, we have described the Rydberg atom array platform, extolling its virtues and attempting to mitigate its weaknesses. We have observed the emergence of random and ergodic behavior from unitary quantum evolution, and have exploited those chaotic effects to benchmark the fidelity of creating highly entangled states of up to 60 atoms. Ultimately, we have tested our system against state-of-the-art quantum and classical processors, finding it is competitive with both, and is entering a regime of computational advantage through quantum simulation.

Now, in the remainder we ask what is next, both for the Rydberg atom array platform, and for the physics we have studied more generally. To tackle this unbridled task, we discuss future work along three axes¹: 1) improving system fidelities, 2) increasing system sizes, 3) exploring new physics.

11.1 IMPROVING FIDELITIES

As we have extensively studied, it is of paramount importance to understand the errors affecting quantum system, and to mitigate them. Not only do such studies inform our approach to the era of error-correction, they also allow us to make more firm conclusions as to the signals obtained on near term devices. To this end, we have spent considerable effort building a sophisticated error model (Chapter 2) which has been cross-checked extensively versus experimental data (Chapters 3, 6, and 8). Using this model, we can explore what are the major limitations along several research directions.

¹Shamelessly adopted/stolen from Ref. [76]

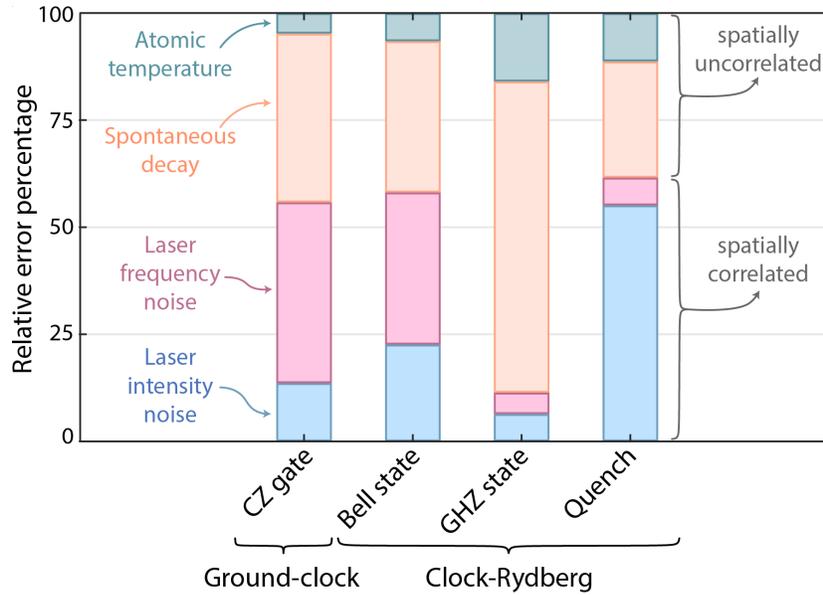


Figure 11.1: **The limitations of quantum operations.** Using our *ab initio* error model (Chapter 2) we determine the major contributing error sources to various quantum operations of interest. Relative error strength is defined as in Fig 8.13, and is $\log(F_i)/\log(\prod F_i)$, where i enumerates the different choices of errors. We see that for different tasks, different error sources become limiting. As such, improvements in fidelity must be carefully tracked to problems of interest.

We study the relative error percentage for the tasks (Fig. 11.1) of: 1) performing CZ-gates with clock-ground qubits (Chapter 4), 2) creating Bell states of clock-Rydberg qubits (Chapter 3), 3) creating clock-Rydberg GHZ states (Chapter 3), or 4) performing many-body quenches ((Chapter 8). Here relative error strength is defined as $\log(F_i)/\log(\prod F_i)$, where i indexes different choices of errors (either laser intensity noise, laser frequency noise, spontaneous decay, or atomic temperature). For all applications, we have checked that the error probability largely factorizes, enabling this metric.

In general, we find that – unfortunately but not surprisingly – different tasks are sensitive to different errors in different proportions. For instance, improving laser intensity noise is paramount for improving simulation of quench dynamics, but largely irrelevant for optimizing GHZ state preparation. Similarly, laser frequency is minimally intrusive for GHZ states and quenches, but is (close to) dominant for performing CZ gates or creating Bell states. While we might of course hope to isotropically improve all aspects of the experiment simultaneously, realistically we need to pick certain noise sources to minimize at a time; this decision must be done in close concert with analysis à la Fig. 11.1. Still, here we broadly lay out steps by which each noise source may be mitigated, each of which is actively being

pursued in our lab. We also note that the use of erasure conversion could change the relative priority of certain upgrades [442].

First, laser intensity noise. Currently, the way in which the laser intensity is stabilized in our system is via a sample-and-hold loop, for which the hold time is a few milliseconds to account for shutter opening times. This means noise at, or faster than, a few kHz is largely invisible to our stabilization. Further, the AOMs we use for stabilization are not fast enough to actively correct for intensity noise *in situ* while pulsing. Thus, we propose two upgrades, both involving EOMs. First, we have bought a noise-eating [443] EOM to be placed directly after the laser output with a tight, always-on feedback loop to eliminate noise close to the MHz timescale which is clearly present in our system (see Fig. 2.5). This noise eater would be followed by our conventional AOM setup for frequency control and further intensity stabilization. Finally, we would introduce another pulse-shaping EOM to act as an ultrafast (nanosecond timescale) optical shutter so as to greatly minimize any servo loop hold times. With these improvements, we expect intensity noise could be substantially reduced.

Next, laser frequency noise. Improvements here could be accomplished through two methods, either 1) improving available Rabi frequencies, or 2) actually minimizing frequency noise. To the first point, most quantum operations are principally sensitive, in general, to noise at their highest energy scale, which typically is set by the drive (Rabi) frequency. We do note that some operations like gates may show some sensitivity at lower frequencies (see some discussion in Ref. [39]), but even there pushing to higher Rabi frequencies generally reduces the relative impact of frequency noise. Intuitively, one can consider that for Rabi frequency, Ω , and detuning Δ , the generalized drive frequency is $\sqrt{\Omega^2 + \Delta^2} \approx \Omega(1 + (\Delta/\Omega)^2/2)$, for $\Omega \gg \Delta$. Thus, frequency noise much lower than Ω is quadratically suppressed. We have purchased and are in the midst of installing an upgraded Rydberg laser which will allow for up to an order of magnitude higher intensity, and thus roughly a $3\times$ increase in Rabi frequency.

Second, we could consider trying to actively reduce the frequency noise. To this end, we note that an obvious and straightforward upgrade to our system is to implement cavity filtering [154], i.e. using the transmission of the high-finesse stabilization cavity directly as a seed for subsequent amplification. This upgrade is underway. Further improvements are possible, but are made difficult by the fact that the laser light is UV, which is generally not amenable to many types of laser stabilization. Still, a passive stabilization approach [444] may allow us to reduce the frequency noise at specific targeted frequencies of interest.

Then, we consider spontaneous decay of the Rydberg state. Notably, we have already largely addressed this issue through active detection of bright state decays via erasure

conversion [26] (Chapter 3), and the same technique could in theory be applied to dark state decay [68]. However, alternatively decay could be minimized either through a (substantial) upgrade of our experimental system to a cryogenic environment [86], or via use of higher lying Rydberg states with improved lifetimes (Fig. 2.8). We note that in this work we have not explored the issue of correlated Rydberg decays in much detail [137]; quantitatively evaluating this process will be paramount for next generation Rydberg atom arrays.

Finally, atomic temperature effects. Again, here we have largely solved such issues through the introduction of our erasure cooling protocol [32] (Chapter 4), via which we achieve nearly ground state cold atoms. Still, for two-dimensional arrays (which we will discuss more in a moment), interaction variations arising from small finite temperatures become significantly more deleterious. Thus, improvements to our overall cooling architecture, such as implementing cavity filtering also on our cooling laser, are advisable.

With these upgrades, it is generally reasonable that about an order of magnitude can be reduced in infidelity across the board for *all* tasks we have considered. Concretely, this would mean CZ gates with fidelities in excess of 0.999, and many-body quenches which are comfortably in the advantage regime (Fig. 10.1).

11.2 INCREASING SYSTEM SIZES

One of the major promised benefits of optical tweezer arrays is their scalability and flexibility. Indeed, numerous experiments with two dimensional arrays [42, 43, 82, 140, 445, 446] of a few hundred atoms have been demonstrated. More recent examples have breached the thousand atom level [83, 84, 141]. In particular, the current state of the art is Ref. [73] (a partner experiment to our own, but one to which the author of this thesis did not directly contribute), where over 6000 atoms are individually trapped in an array of 12000 sites (Fig. 11.2). For our part, we have demonstrated (Chapter 3) an iterative reloading technique for dark-state enhanced loading which enabled dramatically increased scalability (which was then subsequently used for two-dimensional arrays in [141]). Still, as we look to the future it is clear we will need to improve atom numbers beyond the level of a few tens of qubits we currently control. Additionally, for truly large scale quantum computing it will be necessary to link multiple atom array quantum processors together via cavity [139, 204, 447, 448] interconnects [449–451] (Fig. 11.3).

The benefits of larger quantum systems are multiple and manifold. For quantum metrology, clock stability is directly tied to atom number [198, 200]. For quantum simulation, realizing larger systems allows for isolating ‘bulk’ vs ‘edge’ physics [452], with specific applications,

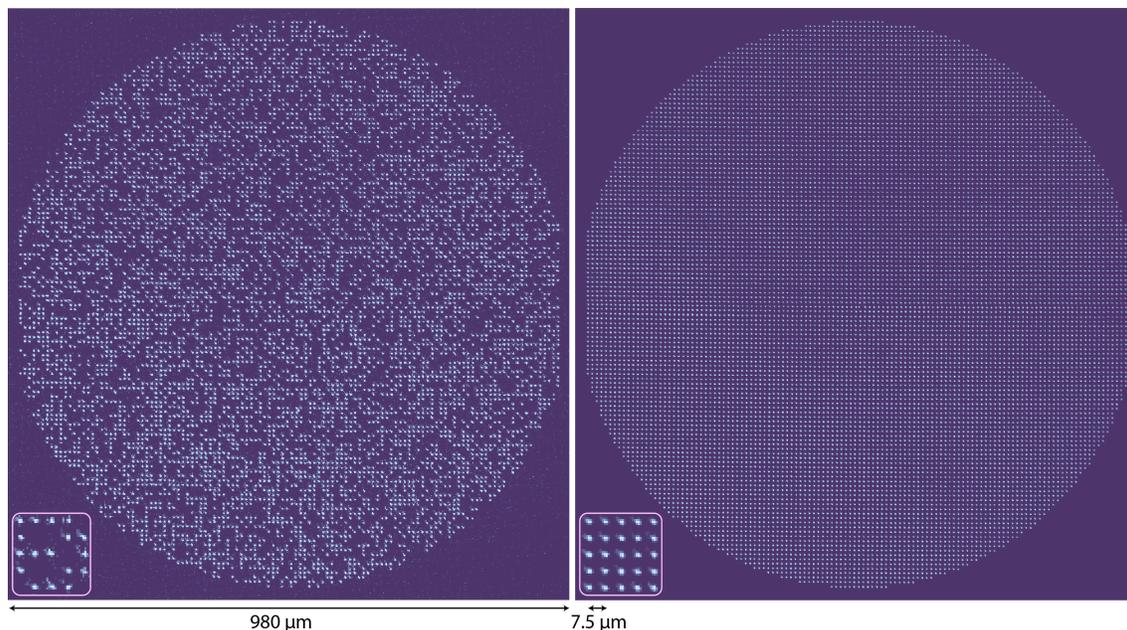


Figure 11.2: **Scaling to optical tweezer arrays with thousands of atoms.** Single-shot (left) and average (right) fluorescence images of a 12000 site optical tweezer array, filled on average with over 6000 individual cesium atoms. This system was demonstrated in Ref. [73] (a partner experiment to our own, though one to which the author of this thesis did not directly contribute), and despite the dramatic improvement in atom number compared to prior works, it also set records for single atom metrics like survival, imaging fidelity, and qubit coherence.

e.g. simulating amorphous quantum systems [453]. Moreover, quantum error correction relies on an extensive number of qubits to minimize errors [50, 454], with schemes likely requiring thousands of physical qubits to achieve roughly a hundred logical qubits [455].

Unfortunately, scaling strontium atom arrays is not easy, given that we typically rely on the magic trapping condition with a tweezer wavelength at 813 nm in order to reliably drive the clock transition. The polarizability of strontium is relatively weak at this wavelength, and further it, as of now, can only be produced at the level of a few watts from conventional lasers [456]. While alternative wavelengths with higher polarizability and greater available laser power could be used, they break the magic condition and come with drawbacks; for instance, it was long thought that 520 nm could be a ‘good’ trapping wavelength for strontium, but it was found to lead to photoionization of the 3P manifold [141]. One approach to achieve higher atom numbers was demonstrated quite successfully in, for instance, Refs. [434, 457] (and earlier works) via the use of a hybrid between a tweezer array and optical lattice. There, the tweezer array utilizes higher polarizability trapping wavelengths to hold more atoms, before transferring them to the power-efficient magic

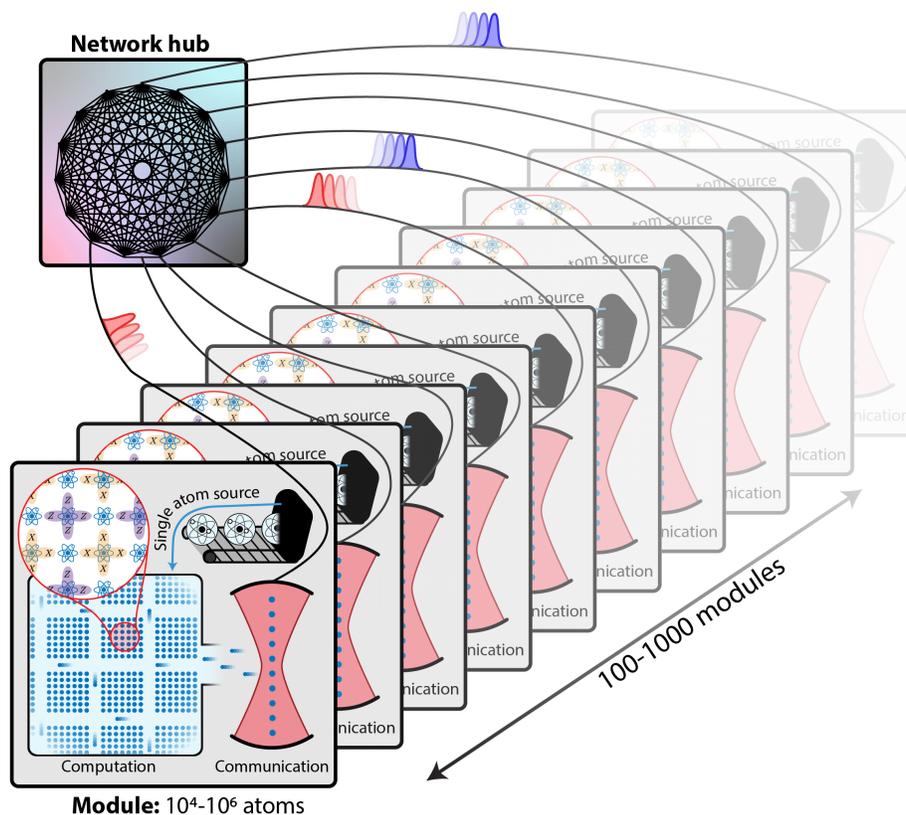


Figure 11.3: **Scaling optical tweezer arrays with cavity integrations.** The most long-term way to scale the sizes of optical tweezer arrays is to modularize them, where individual computing modules are connected via cavity [139, 204, 447, 448] interconnects [449–451].

wavelength lattice.

We plan to scale to relatively conservative system sizes of a few hundred atoms in two-dimensional arrays via a brute force approach: more laser power at 813 nm. We have¹ acquired a coherently combined Ti:Sapph laser system which achieves up to 16 W of output power² for this purpose.

11.2.1 Designing for advantage

One of the most immediate applications of this upgrade will be to repeat similar quantum advantage tests as we have performed in Chapter 10, but now for two-dimensional arrays. The main motivation is that most classical simulation methods largely fail for simulating dynamics of high entanglement entropy states in such systems [458–462]. Further, for such

¹Finally acquired, we should say, given that it was supposed to arrive in the third year of this PhD.

²For how long, who knows.

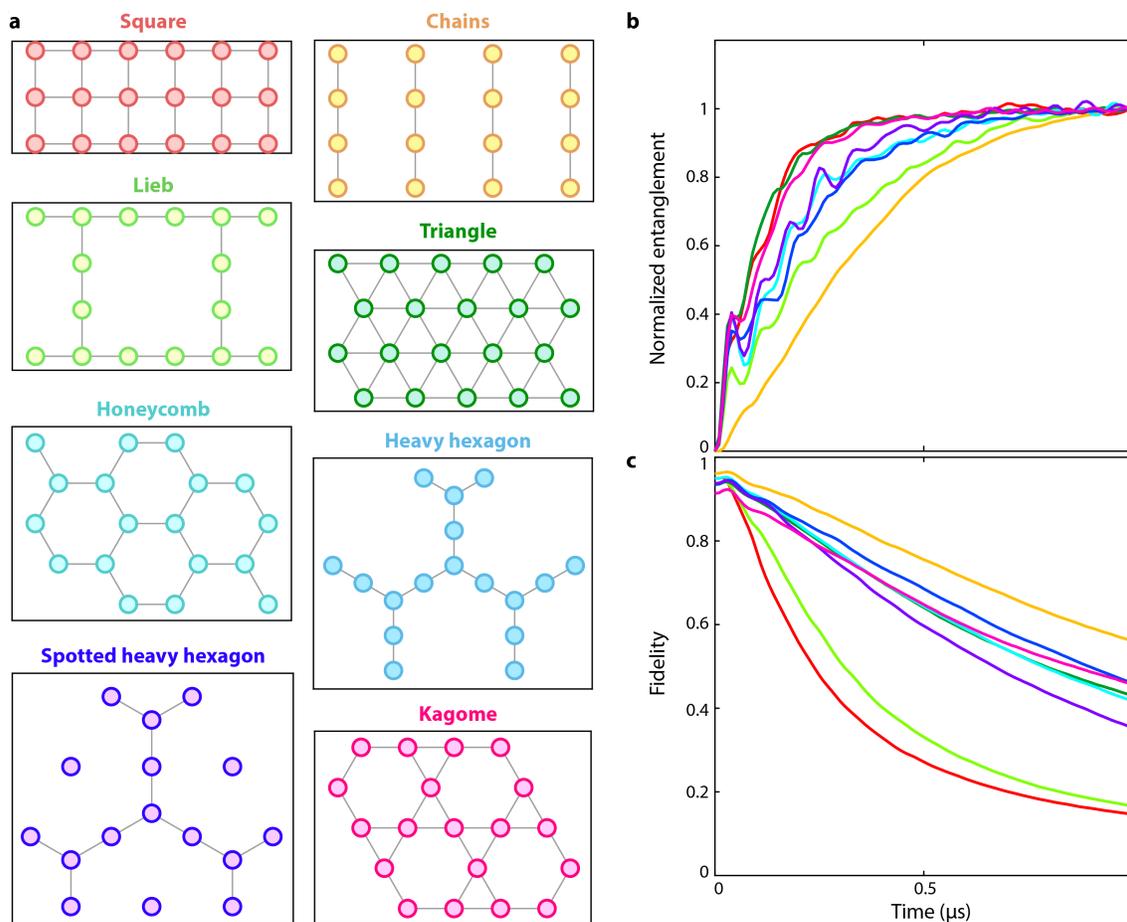


Figure 11.4: **Errors in two-dimensional arrays.** **a**, Unlike one-dimensional systems, two-dimensional arrays feature a number of potential array geometries, each with their own connectivity introducing unique advantages and disadvantages. When considering scaling our fidelity estimation tests to two-dimensional arrays, we perform error model simulations of the eight candidate geometries shown here, where light gray lines represent a blockading interaction between atoms. Note that while some geometries have a higher atom density, they also incur more Rydberg blockade instances, making all the effective Hilbert space dimensions approximately equal. **b**, Entanglement growth of the candidate models in **a**, following resonant quench dynamics. In a two-dimensional it is unclear what the best bipartition of the system is to calculate entanglement; here we perform all possible bipartitions, and choose the one with the *slowest* growth, which we reason is most accurate to the computational complexity of classical simulation. **c**, Fidelity decay from our standard error model. An obstacle to scaling to two-dimensional arrays is the introduction of next-nearest neighbor interactions on the order of the Rabi frequency, for which thermal fluctuations become quite deleterious. Note that here the ‘chains’ array has had the inter-chain distance pre-optimized to maximize fidelity.

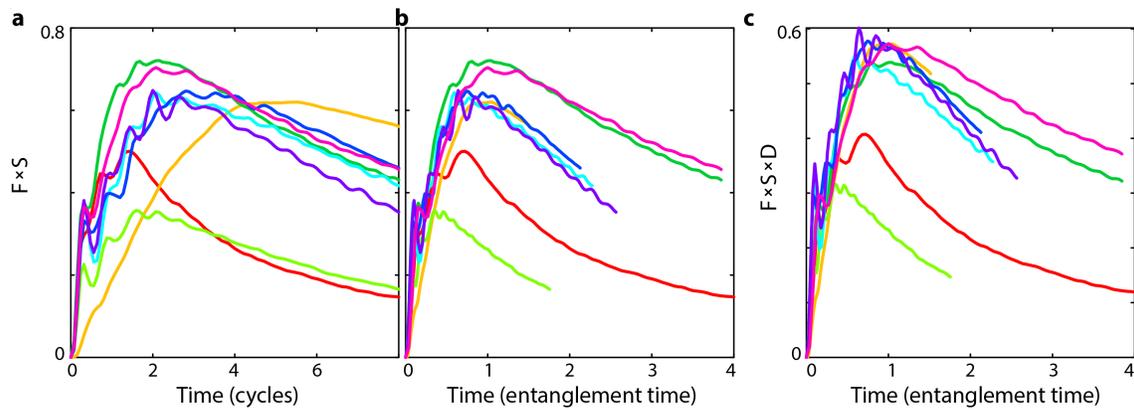


Figure 11.5: **Finding the optimal two-dimensional geometry.** **a**, Entanglement (S) and fidelity (F) product for the candidate geometries in Fig. 11.4a in terms of the absolute evolution time in cycles. $F \times S$ is not a perfect analogue of the mixed state entanglement, but serves as a computationally efficient estimator. **b**, Rescaling each model's time-axis by its unit entanglement time shows many models collapse, but with the highest fidelity-entanglement product for the Kagome lattice and the triangular lattice. **c**, Multiplying the fidelity-entanglement product by the blockade density (D), defined as the logarithm of the blockaded Hilbert space divided by the number of actual atoms, shows many models collapse, without a clear indication of a best model. Thus, it will be an area of active research to find candidate models in two-dimensional systems which bear interesting physics, while reaching computationally-complex regimes.

systems there are complexity-theoretic tests of quantum advantage [386, 463].

However, it is important to take stock of exactly which array geometry will be used in such protocols, in order to maximize fidelity. In Fig 11.4 we consider several such geometries, simulated with our *ab initio* error model. For each we calculate the entanglement entropy growth (normalized by the long time value for fair comparison) and also the fidelity of quench evolution. Note that entanglement entropy here is calculated over all possible system bipartitions, and the final curve is the one with the *slowest growth*, as this is the best analogue of the classical complexity for tensor networks to simulate such systems.

We generally find that there is a trade-off between the entanglement growth rate, and the fidelity reduction of various geometries. In particular, we identify the next nearest neighbor interaction as being particularly deleterious; for two-dimensional systems, the next nearest neighbor interaction can be on the order of the Rabi frequency, meaning noise on the interaction strength (either from temperature or from static positional inhomogeneity) can lead to drastic fidelity reductions. Thus, we generally find the highest fidelity systems are those which minimize next nearest neighbor interactions, for instance in triangular-based arrays.

With these numerics in hand, we can then attempt to find the optimal geometry. To do so, we consider the fidelity-entanglement product, a very rough analog for the mixed state entanglement entropy, and thus the simulation complexity (Fig. 11.5). We find that while plotting in terms of cycles shows clearly optimal geometries, rescaling each time-scale individually by the ‘natural’ units of the entanglement saturation time show most geometries collapse, with some preference for the Kagome or triangular lattice.

However, we must also account for the fact that certain geometries require a greater number of physical atoms to achieve the same blockaded Hilbert space size (because the nearest neighbor coordination number changes from one geometry to the next). As such, we consider the fidelity-entanglement-density product (where the density is the effective number of qubits in the blockade subspace, divided by the physical qubits), for which we find all geometries are roughly equal, with a few low performers like the square array, and perhaps the Kagome lattice performing best (Fig. 11.5c).

Thus, we leave it as an open, interesting question to explore which array geometries bring the most fruitful combination of rapid entanglement growth, and slow fidelity decay.

11.3 EXPLORING NEW PHYSICS

Finally, we turn to the question of what new physics there is to explore with our system. With a high-fidelity, highly-programmable, high-entanglement quantum simulator/processor in hand, the applications are manifold, but here we focus on just three: 1) quantum processing for metrology, 2) uncovering critical phenomena, and 3) investigating unconventional thermalization.

11.3.1 Quantum processing for metrology

First, there are immediate and near-term applications of our system in the realm of quantum metrology. As demonstrated in Refs. [28, 39], access to the tools of a universal quantum processor (Chapter 4) enables new applications in metrology, such as schemes to achieve exponentially improved precision over classical sensors [147, 199–201]. In particular, in Ref. [39] we combined the tools previously demonstrated in this thesis with long-range, coherence preserving movements to enable dynamic circuit reconfigurability [37] (Fig. 11.6a). This for instance enabled the creation of GHZ states with up to 4 atoms (Fig. 4.11). Additionally, in Ref. [39] we demonstrated that further improvements in metrology could be achieved via the use of ancilla-based readout [464], for instance using a single ancilla qubit to readout the multi-qubit parity of an entangled sensor state (Fig. 11.6b).

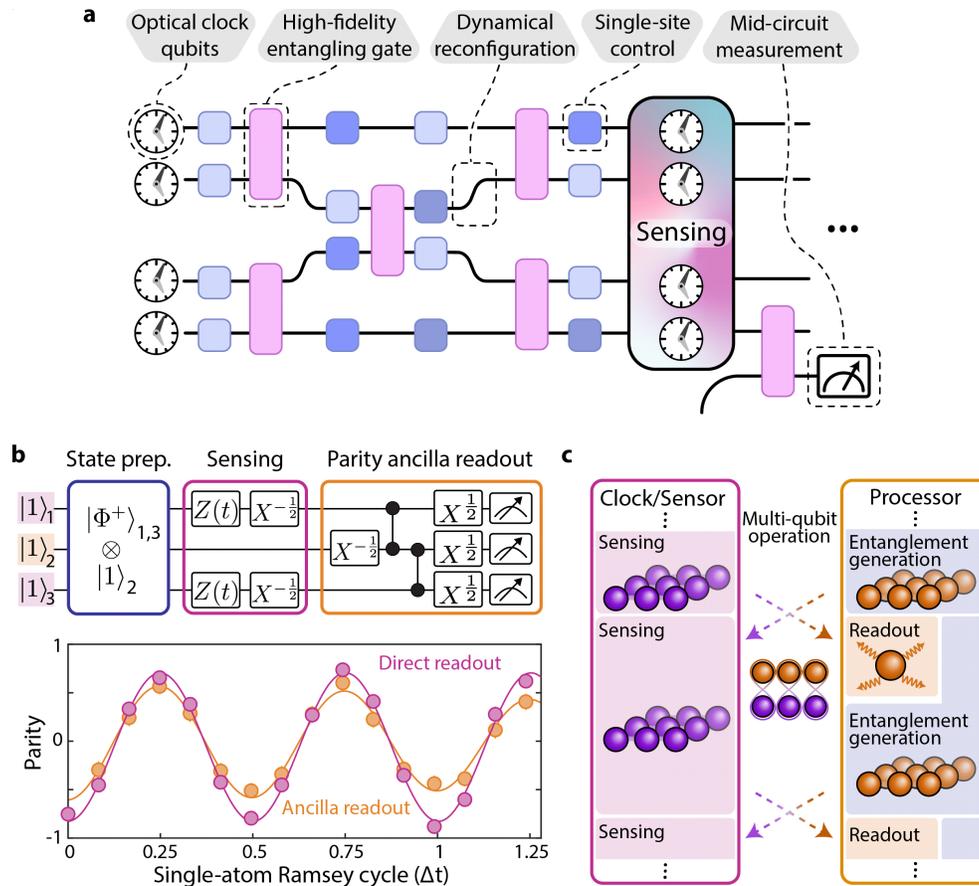


Figure 11.6: **The future of quantum processor-clocks.** **a**, Schematic of the tools of a universal quantum processor mixed with an atomic sensor, all of which we have demonstrated in this thesis (Chapter 4). **b**, Combining these tools together in our platform allows for a myriad of metrologically useful tasks, such as using ancilla qubits to read out the total phase of entangled states to aid sensing. **c**, We envision a future continuous processor-clock hybrid, composed of a dedicated ‘sensing’ module and a ‘processing’ module. While the sensor measures a signal of interest, the processor prepares large scale entangled states to swap with the processor for enhanced metrology, and can further readoff the state of the sensor efficiently through ancilla readout. For more details, see Ref. [39].

With such a capability in hand, there is clear pathway towards a next generation modular clock-processor system (Fig. 11.6c).

In such a system, we imagine two distinct modules which may ideally be composed of two distinct atomic species [140] for ease of control. One module (for instance, strontium atoms), would continuously sense the signal of interest. The other processor module (for instance, rubidium atoms) would serve two purposes. One, while the other module senses the field, the processor would simultaneously prepare an entangled state register. When the sensor has finished a round of measurement, a swap operation [464] would be performed

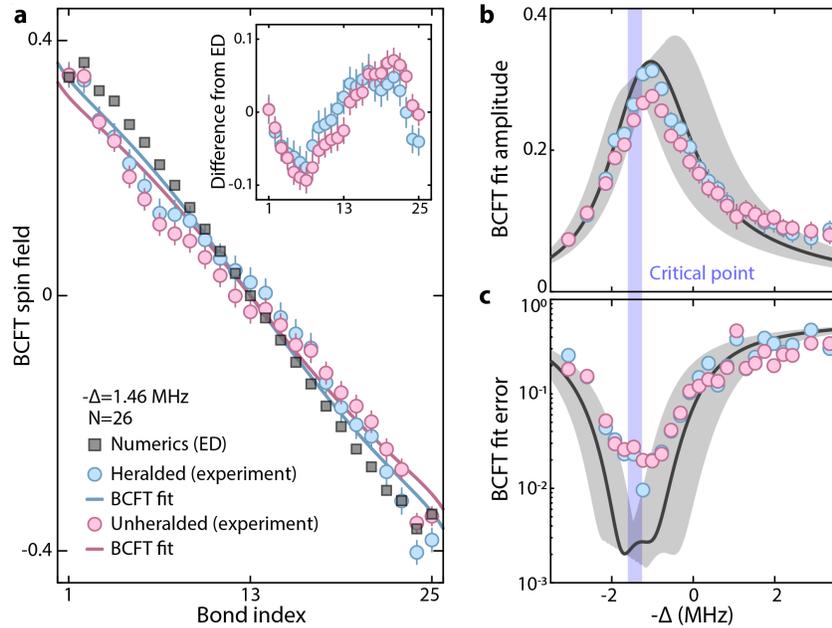


Figure 11.7: **Microscopic conformal field theory physics.** Following Ref. [323], we study boundary conformal field theory (BCFT) observables associated with the phase transition between the disordered phase and the \mathbb{Z}_2 ordered phase. For definitions of all BCFT observables, see Ref. [323]. **a**, Experimental, numerical, and analytic BCFT spin observables for $N = 26$; the agreement between the experiment and the prediction improves slightly with erasure excision. **b**, The BCFT prediction has a free amplitude parameter which is left open in **a**; we find the amplitude is maximized at the predicted critical phase transition point. The phase transition is found via scaling collapse, extrapolated to infinite system size [158]. The shaded area for the numerics comes from uncertainty in our Hamiltonian parameters, in particular the C_6 coefficient. **c**, With the free amplitude, we find the BCFT fit error is minimized at the critical point, and that the fit error is smaller using erasure-excision.

between the sensing and processing modules. Thus, the sensor would be prepared in a fiducial, metrologically-relevant state and could immediately begin sensing once more. Meanwhile, the state of the processor could be readout without crosstalk to the sensor, and the next entangled state could be prepared. This would enable a continuous operation quantum clock achieving Heisenberg scaling with near-zero downtime.

11.3.2 Uncovering critical phenomena

Second, there are fruitful questions of quantum simulation to study with our system. In Chapter 3 we studied spontaneous symmetry breaking across a quantum phase transition, and resolved clear signatures of this process. However, there is further interesting critical phenomena directly at the phase transition itself. Specifically, we follow Ref. [323] and study the *boundary conformal field theory* (BCFT) characterizing the transition between the

disordered and ordered phases in the Rydberg model. We leave most details to Ref. [323], but essentially characterize a single on-site observable, the BCFT spin (akin to the AFM on-site magnetization).

For the 26-atom system we study here, we find good agreement between the experimental BCFT spin profile and that predicted from numerics and also analytically (Fig. 11.7). The BCFT analytic prediction has a free amplitude, which we find is maximized at approximately the critical point, and accordingly we find that fitting the experiment to the BCFT prediction shows a fit error which is minimized at the critical point. Further, the fit error is lessened through the use of erasure excision. Though not a thorough examination of the apparent critical phenomena (in particular, we do not here characterize critical exponents [465]), this suggests an experimental demonstration of microscopic conformal field theory behavior, and indicates the power of erasure-excision to improve such studies more generally.

Still, this investigation is far from complete. The BCFT theory [323] has concrete predictions not just for chains with even length, but also odd length chains. Further, it makes predictions for higher-order observables like correlators, and even string operators related to Majorana modes. All such observables should be visible experimentally, potentially through the use of advanced shadows-based [300, 466] readout methods. Extensions of this study to two-dimensional systems enables studying quantum spin liquids [467]. Then, introducing low-energy dynamics on top of the ground state would enter an unexplored (and classically complex) playground of exciting new physics related to the behavior of (near-)critical systems. Such critical states have information-theoretic properties, such as a cusped Fisher information [159–161]. Interestingly, we believe that tomography of these phenomena may be greatly advanced through the use of speckle-benchmarking techniques, as we have studied extensively with regard to random state ensembles in Chapter 5.

11.3.3 Investigating unconventional thermalization

Finally, there are numerous open questions related to the subject of hydrodynamics [468] and thermalization in our system. In particular, we turn back to the large-scale dataset considered in Chapter 10. At the time, we only studied the global fidelity with which our experiment produced highly entangled target state of interest, but now we investigate the behavior of its low order observables like magnetization and few-body correlators; the system was thermalizing at close to infinite temperature, and so we would not generically expect there to be features of interest.

We focus on the anti-ferromagnetic two-body connected correlator, which for a distance d

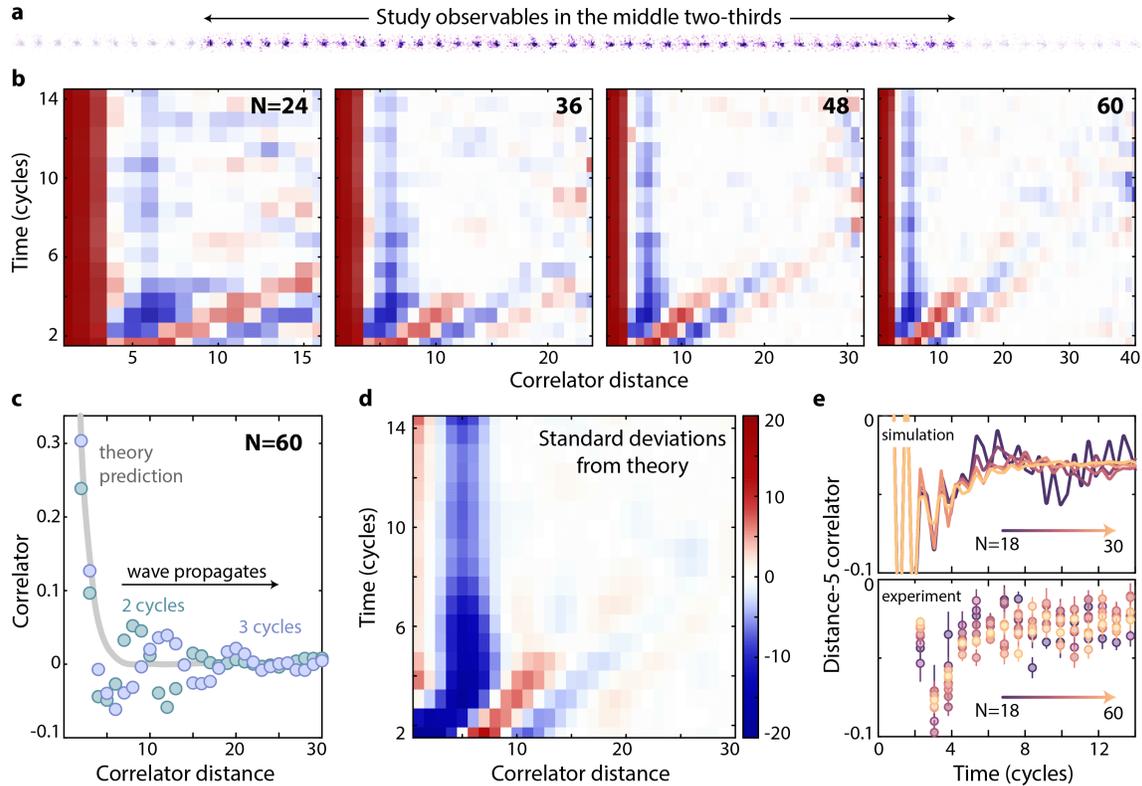


Figure 11.8: Observation of unexpected thermalization. **a**, We study few-body observables isolated to the bulk of large one-dimensional systems following infinite temperature quench dynamics (the same setting as for Chapter 10). **b**, The anti-ferromagnetic connected correlator (see text) show non-trivial behavior consistently across many system sizes, including an apparent ballistic spreading of correlations at early times, and a long-lived, short-range correlator at distance 5 (color scale is arbitrary, but consistent across all plots). **c**, We compare these observations with the expectation of an infinite temperature state in the blockaded manifold, and find neither are consistent with expectations. **d**, We plot the relative error in units of the experimental error bar between the experimental values and the theoretical prediction, and find both the long-lived mode and ballistic propagation are well outside error bars (some temporal smoothing has been applied to the data for visualization). **e**, For both simulations and experiment, the distance-5 correlator saturates to a non-zero value as a function of system-size and time (though more careful simulation shows it very slowly decays towards zero).

between two atomic sites is given by

$$\mathcal{C} = (-1)^d \langle \hat{n}_i \hat{n}_{i+d} \rangle - \langle \hat{n}_i \rangle \langle \hat{n}_{i+d} \rangle, \quad (11.1)$$

where \hat{n}_i is the Rydberg population operator on site i . For a blockaded system thermalized at infinite temperature, we expect this correlator to decay exponentially as

$$\mathcal{C} \approx 0.2 \left(\frac{\varphi - 1}{\varphi} \right)^d, \quad (11.2)$$

where φ is the golden ratio (see Appendix C).

We plot the experimentally found \mathcal{C} for a number of different system sizes, where we isolate the calculation only to the bulk of the one-dimensional chains (Fig. 11.8). We observe clear non-exponential and unexpected behavior, including ballistic spreading of correlations at early times, and a long-lived mode at a distance of 5. These behaviors are clearly distinct from the theoretical expectation, which we typify by plotting

$$r = \frac{\mathcal{C}_{\text{exp}} - \mathcal{C}_{\text{thy}}}{\epsilon_{\text{exp}}} \quad (11.3)$$

where ϵ_{exp} is the error bar on the experimental data points. In other words, the quantity r measures the distance in units of experimental standard deviations between the experiment and theory, which we find to be large both for the long-lived mode and the ballistic correlator (Fig. 11.8d). While analysis of these features is ongoing, we identify similar features to the ballistic spreading in studies of finite temperature hydrodynamics and transport in systems of itinerant particles [469–472]. For the long-lived mode, we believe it is connected to a little studied¹ new family of scar states [474] with significant overlap with the vacuum state. It is interesting and exciting that non-trivial hydrodynamics and non-thermalizing scar states can be found in the setting of a one-dimensional homogeneous spin system under infinite temperature thermalization, and further study is needed to fully characterize these behaviors. Then, with our upgrade to two-dimensional systems, these studies will become even more nontrivial as classical simulations and theory predictions become more difficult.

Even in one-dimension, it is not clear if these phenomena are trivial to study classically. It is conventional wisdom that tensor network based methods are efficient at studying few-body observables, but we can directly test this idea against the results of our high-fidelity experiment. In particular, we plot the anti-ferromagnetic connected correlator for system sizes of $N = 30$ and $N = 60$, which we simulate using our Lightcone-MPS (Section 10.2.4) algorithm; as a reminder this algorithm was found to be by far the most efficient at producing states with the maximum overlap fidelity with highly entangled target states produced by infinite temperature quench dynamics (Chapter 10). Still we find that the two-body correlators do not converge until quite high bond dimension is used in the MPS simulations (Fig. 11.9).

For small system sizes (up to $N = 30$) we can compare the results from low-bond dimension simulations against exact values. Hypothesizing the long-lived distance-5 correlator is the most nontrivial feature to simulate (as it is persistent far beyond the entanglement saturation

¹As far as we can tell, this family is only identified, but not explored, with unsupervised machine learning methods in Ref. [473]

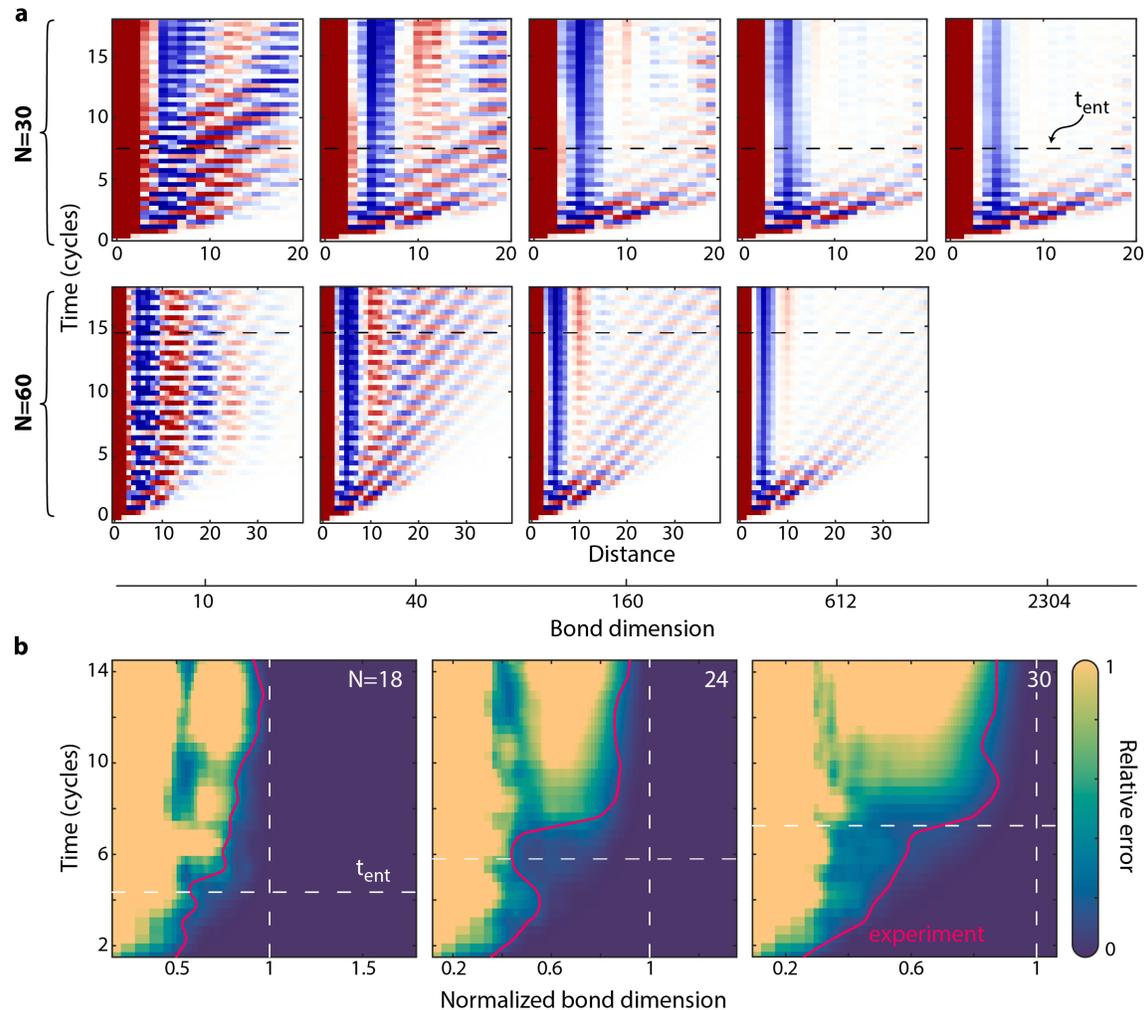


Figure 11.9: Classically complex low-order observables? **a**, We simulate systems sizes $N = 30$ and $N = 60$ using the Lightcone-MPS (Section 10.2.4) algorithm with varying bond dimension, and plot the anti-ferromagnetic connected correlation function. We find correlators continue to evolve as a function of the bond dimension up to the saturation bond dimension (color scale is arbitrary, but consistent across all plots). **b**, For small system sizes, we plot the relative error for the distance-5 correlator between simulations at a low bond dimension to the exact simulations, and find the relative error continually increases as a function of time for limited bond dimensions. Further, we plot the equivalent bond dimension required to match the relative error of the experimentally found correlators (red lines), which show that for long-evolution times nearly the maximum saturation bond dimension is needed to match the experimental accuracy. While we cannot exactly simulate larger systems, we conjecture the trend continues, possibly indicating the experimental observations of low-order observables are the most accurate compared to limited bond dimension classical simulations (though we leave open the possibility that specialized simulation methods may better accomplish this task [475]).

time), we plot the relative error for the distance-5 correlator between the low bond dimension simulations and the exact simulation. We find this relative error continually increases as a function of time, except when the bond dimension surpasses the saturation bond dimension (Fig. 8.17). Further, we study the effective bond dimension needed for the low bond dimension simulations to have a smaller relative error than the experimental data, which we find continually increases; past the entanglement saturation time, roughly the saturation bond dimension is needed to match the experimental accuracy. We are not able to perform exact simulations for system sizes greater than $N = 30$, and so it is not clear if this trend will continue, but we hypothesize it will, perhaps indicating that our high-fidelity experiment might have some advantage at accurately measuring few-body observables already at system sizes of $N \approx 60$.

Of course, this study is not comprehensive, and there may well be classical simulation methods which perform worse than our Lightcone-MPS on the level of global fidelity, but are more accurate for measuring few-body operators [471, 475]. However, it is exciting that we are clearly entering a regime where exact and careful choice of classical algorithm is important to surpass the performance of the experiment even on these ‘simple’ observables. When the experiment is upgraded to a two-dimensional array, such comparisons will become even more nontrivial, and it may soon be the case that the experimental quantum simulator is the most accurate simulacrum for studying phenomena like hydrodynamics and quantum thermalization more generally. It may well be that such tasks hold the key to practical demonstrations of computational advantage through analog quantum simulation [59], a goal well in-line with the overarching aspirations of this thesis.

APPENDIX A

Experimental automated feedback

For many of the datasets considered throughout this work, a large number of samples (many hundreds of thousands of data points) are needed to reach significant conclusions. As the cycle time of the Rydberg atom array is around 1 second (limited mostly by atom loading and imaging), this then entails many days or weeks of data-taking. Even with great effort, it is difficult to passively maintain environmental conditions to such a high-level so as to keep the experimental results temporally consistent over such a long period, so instead we utilize an active feedback approach.

We utilize a home-built control architecture¹ which automatically and periodically interleaves data-taking with feedback to several experimentally relevant parameters. For example, some of these are: 1) the Rydberg transition resonance frequency, 2) the Rydberg Rabi frequency, 3) the clock transition resonance frequency to ensure maximal clock state preparation fidelity, and 4) the Rydberg laser beam alignment. In all cases, rather than using an instrumental signal such as a voltage on a photodiode (which is susceptible to environmental variations like temperature drifts), we instead measure an atomic signal specifically targeted to the observable of interest. After each feedback experiment, the relevant parameters are automatically updated for subsequent measurements (Fig. A.1). Aggregating such measurements, we find that we can control parameters of interest to within a few tenths of a percent over even weeks-long time periods, which is significantly better than the passive stability.

¹Informally named DRAQON (Dynamic and Robust Automated Quantum OperatioN), but that was deemed too ‘cheeky’ by Ivaylo Madjarov so instead we went with ‘MATControl.’ Other potential names under consideration: KRAQEN, EAQLE.

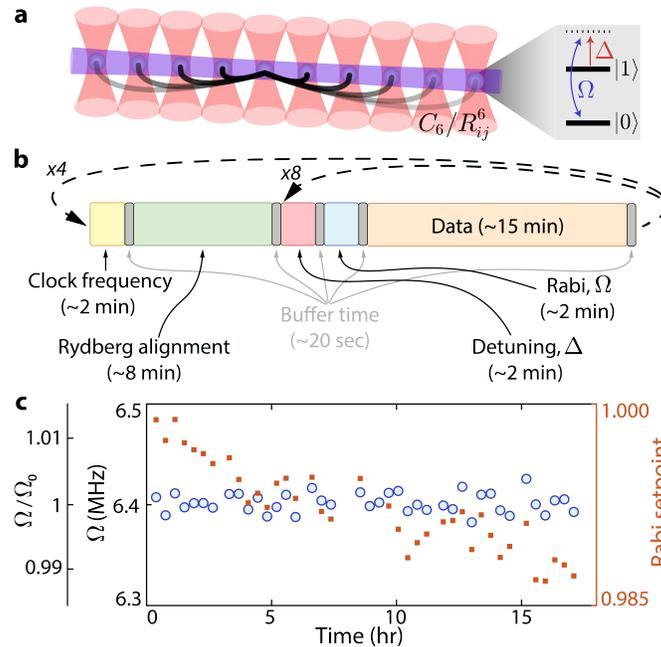


Figure A.1: **Experimental system and parameter feedback.** **a**, Illustration of a Rydberg quantum simulator consisting of strontium-88 atoms trapped in optical tweezers (red funnels). All atoms are driven by a global transverse control field (purple horizontal beam) at a Rabi frequency Ω and a detuning Δ (right panel). The interaction strength is given as C_6/R_{ij}^6 with an interaction constant C_6 and atomic separations R_{ij} between two atoms at site i and j . **b**, Schematic of the experimental feedback scheme. We automatically interleave data taking with feedback to global control parameters and systematic variables through a home-built control architecture; in particular, we feedback to the clock laser frequency (to maintain optimal state preparation fidelity), the Rydberg laser alignment, the Rydberg detuning Δ , and the Rabi frequency Ω . Note that while here we show the Rabi frequency and detuning feedback occurring in two separate blocks, in more recent iterations of the control software we combine them into a single measurement. **c**, Example of the interleaved automatic Rabi frequency stabilization over the course of ≈ 20 hours with no human intervention. Feedback is comprised of performing single-atom Rabi oscillations, fitting the observed Rabi frequency, and updating the laser amplitude, rather than simply stabilizing the laser amplitude against a photodiode reference. While the Rabi frequency setpoint (orange squares) changes over the course of the sequence (due to long-time instabilities like temperature drifts), the measured Rabi frequency (blue circles) stays constant to within $<0.3\%$, with a standard deviation of 0.15% . Similar performance is seen over a matter of weeks, see Fig. 8.4.

Rydberg Rabi frequency and detuning

In order to keep the Rydberg laser's intensity and frequency stable, we measure the spectrum of the laser with Rabi spectroscopy at a fixed, late time (with atoms spaced far apart so as to be non-interacting). Doing so produces many fringes in the spectrum which we then

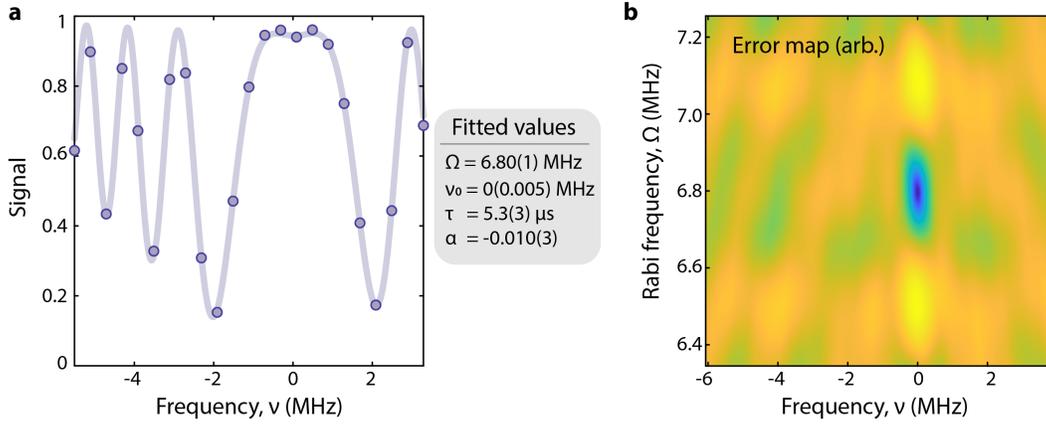


Figure A.2: **Automated feedback on the laser intensity and frequency** **a**, To maintain a constant Rabi frequency and laser detuning, we measure a Rabi spectroscopy signal at a late time (here around $\Omega t = 24\pi$) as a function of detuning; the signal itself is the atomic survival. We then fit the resulting lineshape with Eq. (A.1), which accounts for decoherence sources as well as corrections to the lineshape. Fits are consistent with other ancillary measurements. Note that while the signal around the line center looks relatively flat (which would normally imply insensitivity) the fit is in fact quite precise because of the fringes at larger detunings (note that here for presentational simplicity we have subtracted off ν_0 , but still present its error bar). Indeed, using this fit form, the evolution time does not need to be carefully chosen - the fit will work well regardless. **b**, To make the fitting more robust we perform it in stages (see text), for instance manually computing the least-squares error over choices of Rabi frequency and line center.

fit to simultaneously determine the line center and the Rabi frequency; the later the time, the better the frequency resolution (Fig. A.2). By line center, we mean the programmed frequency on the Rydberg laser such that the laser frequency is resonant with the atomic transition. However, this signal is susceptible to decoherence, but we can include such terms in our fit function. Specifically, we fit the ground state population signal, S , at a time t as a function of the programmed frequency, ν , as

$$\begin{aligned}
 S(\nu) = & 1 - \left(\frac{\Omega^2}{\Omega^2 + (\nu - \nu_0)^2} \right) \times \\
 & \left(\left(1 - \left(\cos(2\sqrt{\Omega^2 + (\nu - \nu_0)^2} \pi t) \right) \left(\exp(-(t/\tau)^2) \right) \right) / 2 \right) \times \\
 & \left(1 + \alpha(\nu - \nu_0) \right) \times \\
 & \left(\exp(-\gamma(\nu - \nu_0)^2) \right), \tag{A.1}
 \end{aligned}$$

with fit parameters $[\Omega, \nu_0, \tau, \gamma, \alpha]$, where Ω is the Rabi frequency, ν_0 is the line center, and τ is an effective Gaussian decay rate. Then the parameter γ is a numerically found factor which accounts for the finite turn-on time of the AOM (Section 2.3.2): because the AOM

does not output a perfect square pulse, the Rabi lineshape (which is related to the Fourier transform of the intensity pulse) is narrowed. The exact value of γ should be numerically found from simulations of measured AOM turn-on dynamics. The final parameter α is a phenomenological one which we introduce to account for the fact that a slight misalignment of the atoms with respect to the laser leads to a time-dependent variation in Rabi frequency as well as laser frequency (Section 2.3.2). In general, we find the value of α is well-correlated with our alignment; if the alignment is good (described below), then α will be consistent with 0.

In practice we find it is most robust to fit make a quick estimate of τ as

$$\tilde{\tau} = \tilde{\Omega} / \sqrt{\log(C)} \quad (\text{A.2})$$

where $\tilde{\Omega}$ is an initial guess for the Rabi frequency, and C is an estimate of the contrast (the difference between the maximum and minimum observed values) in the experimental line shape. We then manually find the optimal choices of Ω and ν_0 (fixing $\tau = \tilde{\tau}$ and $\alpha = 0$) by explicitly rasterizing over a two-dimensional region in their parameter space, calculating the squared error for each point in that space, and finding the corresponding minimum. This error space shows multiple minima spaced at 2π increments, but we find the lowest error minima is always the correct one. We then take the corresponding manually found Ω and ν_0 as initial parameters and refit Eq. (A.1), including with free τ . We find this procedure is quite robust at estimating the Rabi frequency and line center simultaneously in only ≈ 150 experimental shots.

Clock frequency

For feedback on the clock frequency, we utilize two different methods. The first is essentially the same as described previously for feedback on the Rydberg laser. However, our clock laser requires more frequent frequency feedback than the Rydberg laser (for instance, we see clear signatures of noise on ≈ 10 minute timescales that must be corrected for), so we supplement with a quicker and ‘dirtier’ feedback. We first carefully feedback on the frequency as described above, then we perform Ramsey spectroscopy.

The clock self-light shifts itself [69, 76], and so we carefully select an offset frequency during the Ramsey dark time to cancel this effect, such that ideally the Ramsey signal should show no oscillations, and should be quite flat around unity (assuming no decoherence). We then set the Ramsey dark time to a very late time, and sweep the frequency during the dark time in a very narrow range. We fit the result with a Gaussian (or a quadratic in some cases), and use the result to adjust the clock frequency. This method takes fewer shots, but we find it can be more unstable on longer time periods, because it is susceptible to ‘fringe-hopping’

by 2π . Thus our normal clock feedback consists of one round of the normal Rabi-based feedback (which fixes the absolute frequency) for every roughly 10 rounds of Ramsey-based feedback (which efficiently corrects for small relative frequency offsets).

Alignment

Correctly optimizing the alignment between the Rydberg laser and the atoms is crucial: the laser waist-radius is only $\approx 15 \mu\text{m}$, meaning relative displacements by as much as one micron can drastically impact fidelity (see Section 2.3.2). To perform automated feedback of the alignment, we illuminate the atoms with a fixed, long-time pulse, and then raster the Rydberg beam across the array. What signal does this produce? Well, if we tune pulse area to be approximately 1π when aligned, the resultant signal will be a monomodal peak. If instead the pulse area is 2π , the signal will be a bimodal peak. If the pulse area is 1.5π , we get something in between (Fig. A.3).

Assuming the laser beam has a Gaussian mode with waist-radius w_0 , the Rabi frequency, Ω , at a displacement, x , is given by $\Omega \propto \exp(-x^2/w_0^2)$. Thus, the signal we expect to measure is

$$S(x) = \frac{1}{2} \left(1 - \left(\cos \left(\exp \left(-\frac{x^2}{w_0^2} \right) t \times \pi \right) \left(\exp \left(-\frac{x^2}{a^2} \right) \right) \right) \right), \quad (\text{A.3})$$

where a is a fit parameter roughly capturing the effect of increased intensity noise from misalignment (in reality the functional form is slightly more complicated, see Section 2.3.2). In a narrow region around the central peak of the signal we can largely ignore this decoherence term. To calculate the sensitivity of this signal, we consider the second derivative around $x = 0$ at $t = 0.5, 2.5, 4.5, \dots$ (i.e. $\pi/2$ times, chosen such that the integer component is even so the second derivative is positive), for which we find the sensitivity increases linearly with the evolution time. This means in principle we can gain arbitrary statistical certainty by successively increasing the pulse area. Indeed, we are able to see alignment accuracies at the few tens-of-micron level (Fig. A.3b).

Another important aspect of this technique is that when we raster the laser across the atoms, we can resolve the signal on an atom-resolved basis. This allows us to image not just the mean position misalignment of the laser beam, but also misalignment in its angle, and misalignment along the z-position (i.e. along the array axis). This is why in practice we do not actually use Eq. (A.3) to fit the experimental signal, but instead perform a manual sweep across the relevant parameters of beam position, angle, divergence, finding the set which best matches experiment. We can also independently measure the angle for instance by taking a measurement like Fig. A.3b, fitting the first-order troughs or peaks symmetrically outside the central feature, and then plotting their slopes as a function of position along the

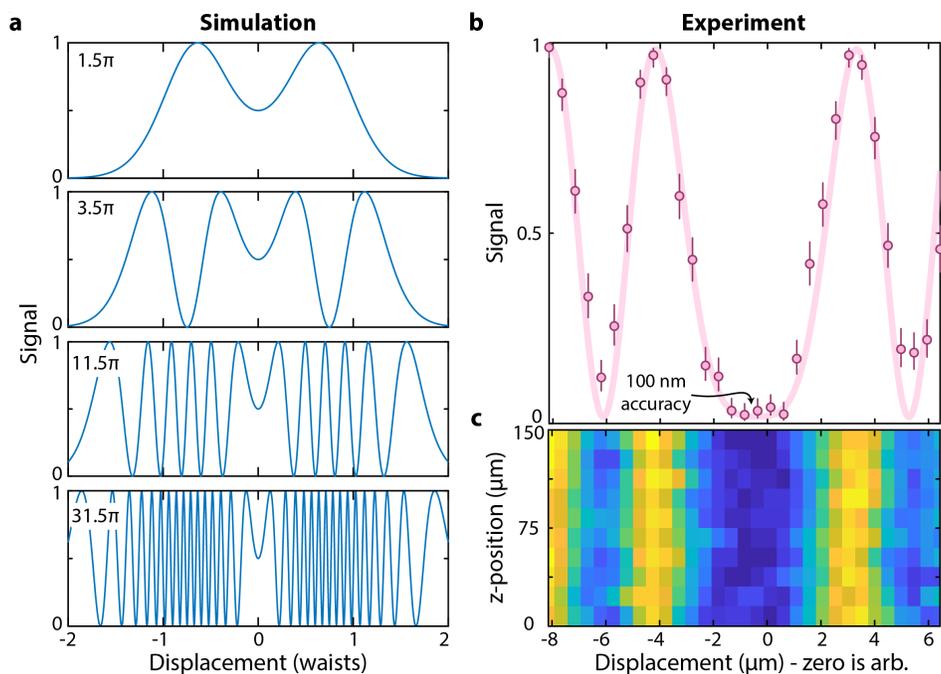


Figure A.3: **Automated feedback on the laser alignment** **a**, For a given aligned pulse area (number in top left of each plot), the signal when rastering the laser beam across the atoms will change, where the sensitivity (i.e. second derivative at zero displacement) of the signal increases linearly with the pulse area. Here we show the case with no decoherence from misalignment, though note that such effects are quadratically suppressed in the narrow region around the optimum. **b**, For a large pulse area (here $\approx 30\pi$ the resultant raster can be quite accurate down to the level of $\lesssim 100$ nm. **c**, When rastering the beam across the atom array, we can take a full three-dimensional map of the beam intensity (similar to work in cavities [139]). Here we show the case of a well aligned signal, but if the beam had an angle with respect to the atoms, or was displaced in the z-direction (the array axis), it would be clearly visible, and correctable.

tweezer array axis. For typical alignments, we perform first the alignment in the vertical direction and then the horizontal, each taking roughly 200 experimental shots.

APPENDIX B

Numerical simulation methods

Here we give a (very) brief tutorial on classically simulating quantum systems. While this topic is ubiquitous in the field – and covered by several simulation packages [476, 477] – it is surprisingly difficult to find a good, tutorial-esque introduction to these concepts. A notable exception is Ref. [478], which is a clear, comprehensive reference to consult with. Further, Refs. [479, 480] are good tutorials for specific aspects of classical simulation with explicit examples and code for MATLAB and Mathematica. The goal of this chapter is not to reproduce all the results of these established texts, but simply to give a practical introduction in a way which would have been helpful early in this author’s PhD process.

B.1 HAMILTONIANS AND SYMMETRIES

To start, we will consider one of the simplest problems of interest: simulation of on-resonance two-qubit Rydberg dynamics. Two qubits are easily simulable on any classical computer, but we use this model system to introduce tools to efficiently represent and simulate quantum systems applicable to larger systems.

The full Hamiltonian is given in Eq. (2.1), but here we will work with the simplified Hamiltonian (with $\hbar = 1$ throughout)¹

$$\hat{H} = \frac{\Omega}{2}(\hat{X} \otimes \hat{I} + \hat{I} \otimes \hat{X}) + V(\hat{n} \otimes \hat{n}), \quad (\text{B.1})$$

¹While basic, it is worth stressing that when people generally write sums of operators, like $\hat{H} = \sum \hat{X}$, they are actually performing a *Kronecker sum*, and inserting identity operations for unaffected subspaces. For example $\hat{H} = \sum_{i=1}^3 \hat{X}_i = \hat{X} \otimes \hat{I} \otimes \hat{I} + \hat{I} \otimes \hat{X} \otimes \hat{I} + \hat{I} \otimes \hat{I} \otimes \hat{X}$.

where

$$\hat{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \hat{I} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \hat{n} = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}. \quad (\text{B.2})$$

The most straightforward way to construct the Hamiltonian is simply to directly perform Kronecker products on the underlying matrices, yielding a matrix, \hat{H} , operating on the standard computational basis, \vec{b}

$$\hat{H} = \begin{bmatrix} 0 & \frac{\Omega}{2} & \frac{\Omega}{2} & 0 \\ \frac{\Omega}{2} & 0 & 0 & \frac{\Omega}{2} \\ \frac{\Omega}{2} & 0 & 0 & \frac{\Omega}{2} \\ 0 & \frac{\Omega}{2} & \frac{\Omega}{2} & V \end{bmatrix}, \quad \vec{b} = \begin{bmatrix} 00 \\ 01 \\ 10 \\ 11 \end{bmatrix}. \quad (\text{B.3})$$

To perform time-evolution under this Hamiltonian, we form the unitary $\hat{U} = \exp -i\hat{H}t$, which then can be multiplied onto a vector representation of some arbitrary state, $|\psi\rangle$. Explicitly, this operation requires us to diagonalize the matrix representation of \hat{H} , and as such is known as *exact diagonalization* (ED). While it is – as the name implies – exact, this method is also computationally inefficient. Matrix diagonalization is an expensive operation, scaling as $\mathcal{O}(m^3)$ where m is the linear matrix dimension; in our case $m = 2^N$, where N is the number of atoms, the cost of ED goes as $\mathcal{O}(8^N)$, and thus quickly becomes computationally intractable for system sizes of $N \gtrsim 13$ without much effort (or patience).

To accommodate this problem, we need to either reduce the matrix size, or use a different method of time evolution which does not require explicit diagonalization. To start, we focus on the first point, and ask how we can simplify the Hamiltonian using our physical intuition for the problem. In particular, we can use physical symmetries inherent in \hat{H} to reduce it to a simpler basis. Explicitly, in the case of the two-qubit system considered here, we see there is a swap-symmetry (or *reflection symmetry* more generally) between the two qubits. We can explicitly write the matrix representation of this symmetry operation (dubbed the *transformation matrix*),

$$\hat{P} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (\text{B.4})$$

which we can explicitly verify commutes with \hat{H} , i.e. $[\hat{H}, \hat{P}] = 0$. We want to rewrite \hat{H} accounting for the symmetry operation of \hat{P} , so first diagonalize¹ \hat{P} , and construct the

¹I know, our goal here was to get away from the problem of diagonalization, and yet here we do it again. In a moment we will introduce a more efficient method for this process which does not require explicit diagonalization of \hat{P} .

matrix \hat{R} , where the columns are eigenvectors of \hat{P} :

$$\hat{R} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ \frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} & 0 \\ -\frac{1}{\sqrt{2}} & 0 & \frac{1}{\sqrt{2}} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (\text{B.5})$$

This is the change-of-basis matrix from the eigenbasis of \hat{P} to the standard basis, \vec{b} . The order of the columns in \hat{R} is technically arbitrary, but should be done at the very least such that eigenvectors with equal eigenvalues should be placed next to each other (in a ‘block’).

In the following, let us denote the eigenbasis of \hat{P} with a prime-designation (e.g. $|\psi'\rangle$), and the lack of a prime indicates it is written in the standard basis, \vec{b} . We know that for any such basis transformation where $\hat{R}|\psi'\rangle = |\psi\rangle$, expectation values stay the same: $\langle\psi|\hat{H}|\psi\rangle = \langle\psi'|\hat{H}'|\psi'\rangle = \langle\psi|\hat{R}\hat{H}'\hat{R}^{-1}|\psi\rangle$, and thus

$$\hat{H}' = \hat{R}^{-1}\hat{H}\hat{R} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{\Omega}{\sqrt{2}} & 0 \\ 0 & \frac{\Omega}{\sqrt{2}} & 0 & \frac{\Omega}{\sqrt{2}} \\ 0 & 0 & \frac{\Omega}{\sqrt{2}} & V \end{bmatrix}, \quad \vec{b}' = \begin{bmatrix} 01 - 10 \\ 00 \\ 10 + 01 \\ 11 \end{bmatrix}, \quad (\text{B.6})$$

where \vec{b}' is the new basis.

Importantly \hat{H}' takes on a block-diagonal form, where the first column/row corresponds to the antisymmetric basis $\{(|01\rangle - |10\rangle)/\sqrt{2}\}$, and the other columns/rows correspond to the symmetric basis $\{|00\rangle, (|01\rangle + |10\rangle)/\sqrt{2}, |11\rangle\}$. This makes sense – the swap transformation matrix has simultaneous positive eigenvalues for symmetric states, and negative eigenvalue for antisymmetric states.

Now, the key realization is that if our initial state lives in only one of these subspaces (e.g. our initial state is $|\psi_0\rangle = |00\rangle$) then its dynamics will be entirely constrained to the symmetric subspace. This means the effective Hilbert space is only 3-dimensional, rather than 4-dimensional, and thus any corresponding time-evolution can be performed more efficiently. If we now assume the Rydberg blockade condition (i.e. $V \gg \Omega$), then we can further reduce the to a two-dimensional subspace¹.

This principle is quite general: for any Hamiltonian we wish to simulate, we can always work in a reduced state space by identifying symmetries of the Hamiltonian, block-diagonalizing

¹Astute readers will also note that in this reduced subspace the Rabi frequency is effectively increased by a factor of $\sqrt{2}$, as is well-known for blockaded Rabi oscillations [34].

and proceeding in the lower dimensional manifold. Examples of symmetries commonly exploited in this way are: *translation symmetry*, often found in systems with periodic boundary conditions, *charge conservation*, where the total number of excitations or particles is invariant, *reflection symmetry*, the higher-order analog of the swap symmetry we have already considered, and more. In all cases, write down the transformation matrix, diagonalize it to find the change of basis matrix, and then block diagonalize the Hamiltonian accordingly.

While nice in principle we have not really solved the issue: our method still explicitly requires us to diagonalize a large matrix. Thus, the proper way (which is unfortunately much more cumbersome) is to eschew diagonalization altogether, and instead to construct the Hamiltonian efficiently term-by-term using bitwise operations. This method also turns out to be far more efficient for Hamiltonian creation than simply repeatedly applying Kronecker products. We cannot give a better explanation of this process than the one found in Chapter 4 of Ref. [478], and so direct readers there once they are comfortable with the basic outline we have presented above. Further references are for instance Refs. [481, 482].

B.2 TIME-EVOLUTION

With the Hamiltonian in hand, one needs to actually perform time evolution with it. In its reduced form, it might be possible to perform exact diagonalization efficiently. Indeed, for instance for systems with translation symmetry the reduced Hilbert space can be substantially smaller than the naïve one, and so Hamiltonians of over 30 qubits can be diagonalized. However, this is not a general solution to the problem. There are a sea of methods for performing approximate time-evolution in the literature, most mentioned in one way or another in Ref. [476], but also many are covered in Chapter 10 of this thesis. Many of these methods are tensor network based for which good general tutorials are Refs. [392, 483]. Here we briefly describe just a few of these ideas.

Krylov evolution

One popular non-tensor network based approach for efficient Hamiltonian simulation is known as the Krylov subspace method, or Lanczos method more broadly in science and engineering disciplines [478]. Essentially, the idea underlying Krylov methods is that the chief difficulty in exact diagonalization is the explicitly finding the matrix exponential $\hat{U} = \exp(-i\hat{H}t)$. However, we do not actually care about \hat{U} , but instead only care about its action on a state of interest, $|\psi_0\rangle$. Thus, we are always going to be concerned with solving problems of the form $|\psi_1\rangle = \hat{U}|\psi_0\rangle$. Heuristically, and blithely throwing away lots of constants, we can essentially think of this operation as $|\psi_1\rangle = (1 + \hat{H} + \hat{H}^2 + \hat{H}^3 \dots)|\psi_0\rangle$. Krylov

methods then take advantage of this expansion by writing basis vectors, $|\zeta_0\rangle = |\psi_0\rangle$, $|\zeta_1\rangle = \hat{H}^1|\psi_0\rangle$, $|\zeta_2\rangle = \hat{H}^2|\psi_0\rangle$, $|\zeta_3\rangle = \hat{H}^3|\psi_0\rangle \dots$, and then performing the corresponding unitary rotation only in this basis. Thus, it is effectively approximating $|\psi_1\rangle$ in the span of the $|\zeta\rangle$ vectors. The size of the basis of course depends upon the problem of interest, but for spin systems, a basis size of a few tens to a hundred generally suffices. While certainly possible to write one's own Krylov method implementation, they are quite well studied with efficient libraries available for almost all languages. The key thing to remember is that ultimately they are essentially just a Taylor expansion of the matrix exponential.

Circuit evolution

Often in analog quantum simulation we are given an initial state, and a global Hamiltonian affecting all atoms continuously. However, for digital quantum processors executing quantum circuits the task is quite different. In particular, for circuits we generally assume the quantum evolution occurs only on a small (e.g. one or two) number of qubits independently at a time. How does this help us perform more efficient time-evolution? Consider a state of e.g. $N = 5$ qubits, initially in the vacuum state $|\psi_0\rangle = |0\rangle^{\otimes N}$. Let's say we want to apply a round of randomized single qubit gates, $\hat{O}^{(i)}$, to every qubit, where the superscript denotes the qubit index. To do so we could explicitly construct the global unitary as

$$\hat{U}_1 = \sum_i \hat{O}^{(i)} \quad (\text{B.7})$$

which we then apply to $|\psi\rangle$. However, doing the required Kronecker products is costly, and, as it turns out, unnecessary via use of clever matrix reshaping tricks.

Consider a concrete representation of $|\psi\rangle$, a 32×1 vector. While we could multiply this vector by a 32×32 unitary matrix, we can instead reshape it in order to apply evolution on just a subset of qubits at a time. For instance, to apply the operation $\hat{O}^{(3)}$ alone to the third qubit, we do the following.

1. Reshape the 32×1 state vector into a $2 \times 2 \times 2 \times 2 \times 2$ tensor.
2. Permute the tensor indices so the ordering transforms as $(1, 2, 3, 4, 5) \rightarrow (3, 1, 2, 4, 5)$.
3. Reshape the tensor into a 2×16 matrix.
4. Left multiply by the 2×2 operator $\hat{O}^{(3)}$.
5. Reshape into a $2 \times 2 \times 2 \times 2 \times 2$ tensor, and undo the permutation in (2).
6. Reshape back into a 32×1 state vector.

This same procedure applies to multi-qubit gates, simply with judicious choice of reshaping and permutation.

Reshaping tricks like this are quite useful for a number of quantum operations. For instance, let's say one wants to obtain the reduced density matrix on just qubits 1, 3 and 5 in the above example. Then the same basic procedure applies

1. Reshape the 32×1 state vector into a $2 \times 2 \times 2 \times 2 \times 2$ tensor.
2. Permute the tensor indices as $(1, 2, 3, 4, 5) \rightarrow (1, 3, 5, 2, 4)$.
3. Reshape the tensor into a 8×4 matrix.
4. Right multiply by the matrix transpose, yielding the 8×8 reduced density matrix.

With straightforward changes to the above basic procedure, one can also measure local observables etc. In general, gaining familiarity and confidence with matrix manipulations like reshaping and permutations will be greatly beneficial for practical classical simulation.

Hamiltonian circuits

Finally, we discuss how to import some of the best features of efficient circuit simulation into the realm of analog quantum simulation. This discussion is heavily based on the theoretical work of Ref. [396]. Section 10.2.4 already contains an introduction to this topic, but in the context of matrix product states (MPS); here we will have a less formal discussion using exact diagonalization rather than MPS, with the goal being that readers can use this introduction as a base to learn more details in the aforementioned sources. As a brief note, our discussion here will be on a slightly less optimal, but easier to understand, version of the algorithm in Ref. [396], while Section 10.2.4 describes a more optimal implementation.

As we have just seen, it can be more efficient to break a large quantum evolution up into a circuit model in order to apply quantum operations without the overhead of taking large scale Kronecker products and matrix diagonalizations. It would be ideal if the same could be performed for analog quantum simulations. It is straightforward to break a continuous-time Hamiltonian up into discrete time steps (i.e. Trotterization), but it is less clear how to break a continuous-space Hamiltonian (i.e. one with long-distance interactions) into discrete chunks. Ref. [396] proposes to solve this problem in a clever way through a combination of forward time, and backward time evolution (Fig. B.1).

Let's say we have a system of $N = 18$ interacting qubits in a one-dimensional chain, where the interactions go beyond nearest neighbor, but are not too long-ranged (e.g. Rydberg interactions). The resulting unitary evolution operator is $2^{18} \times 2^{18}$, which is clearly far too large for exact diagonalization, and is costly even for Krylov methods. However, here we must draw on our physical intuition. While interactions are long-ranged, entanglement across the one-dimensional chain cannot spread infinitely quickly, and is instead limited by

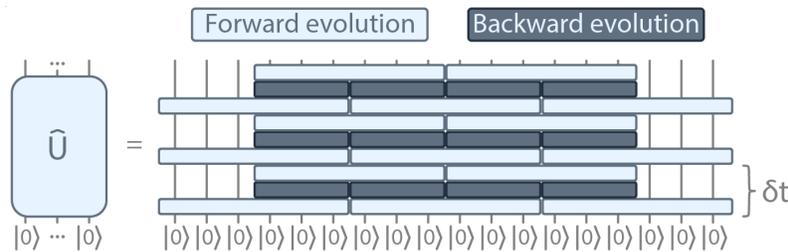


Figure B.1: **The Lightcone unitary decomposition** As theoretically proposed in Ref. [396], and explained in greater detail both in this appendix and in Section 10.2.4, we perform efficient Hamiltonian unitary evolution by exploiting the relatively short-ranged interaction. Specifically, we both Trotterize the unitary evolution in time, and also in space. We use a combination of forward and backwards evolution in order to keep Trotterization commutator errors to an optimal minimum.

the Lieb-Robinson velocity [319]. This means that if we just evolve for a short amount of time, δt , atoms are not able to ‘feel’ the effects of atoms further than few sites away.

To exploit this behavior, we break our chain of 18 atoms into a group, \mathcal{G}_1 , of three consecutive blocks of six atoms each (i.e. block one has atoms 1-6, block 2 has atoms 7-12, block 3 has atoms 13-18) (see Fig. B.1, though note it contains additional complexity we will address shortly). For each, we construct the Hamiltonian *only* in the block, with no external coupling. These Hamiltonians have a linear dimension of only 64, and are easy to diagonalize to produce a block-unitary operator. We can then apply these unitaries to the three chunks separately, and time-evolve the system forward efficiently. Then, for the next time step, we could for instance form a group, \mathcal{G}_3 , of two blocks spanning the interface (i.e. block one with atoms 4-8, and block two with atoms 9-15). We could again efficiently construct their time-evolution operators, and then on alternating time-steps going forward we alternate unitary evolution on \mathcal{G}_1 or \mathcal{G}_3 . For small enough δt , the result will accurately track the ideal result of exact diagonalization over the whole system.

There are some caveats and improvements to be made. First (and less meaningful), in that implementation we are assuming that a block of six atoms is large enough to capture the long-ranged Hamiltonian accurately. Why not choose a block of four atoms, or two? Ultimately, the size of the block simply comes down to the Lieb-Robinson velocity - for a smaller block size, the δt must be accordingly lowered to maintain accuracy. It is in general a good idea to set the block size as big as it needs to be, but no bigger, which for the Rydberg Hamiltonian we find (Section 10.2.4) to be a block size of around 6.

Second (and more meaningfully), in the algorithm above, carefully keeping track of commutators shows a growing error due to the alternation between \mathcal{G}_1 and \mathcal{G}_3 . However, Ref. [396]

proposes a solution. We introduce a new group, \mathcal{G}_2 , of blocks of just three atoms (block one with atoms 4-6, block two with atoms 7-9, block three with atoms 10-12, block four with atoms 13-15) – see Fig. B.1. The block in \mathcal{G}_2 straddle the interface between blocks in \mathcal{G}_1 . The claim of Ref. [396] is that errors are optimally reduced if the time-evolution is done as follows:

1. Perform forward time evolution on blocks of atoms in \mathcal{G}_1 .
2. Perform *backwards* time evolution on blocks of atoms in \mathcal{G}_2 .
3. Perform forward time evolution on blocks of atoms in \mathcal{G}_3 .

The reader may check for themselves that to first order each individual atom experiences the same total evolution, but the introduction of backwards evolution next to the block interfaces eliminates higher-order terms in the Trotterization commutator expansion, and thus substantially reduces such Trotterization errors.

Crucially, for time-independent evolution the unitary operators in the different groups can be precomputed and stored in memory (as they are relatively small), and then applied ad nauseum throughout the evolution. Further optimizations are possible, for instance by combining groups \mathcal{G}_1 and \mathcal{G}_2 during the precomputation step. Additionally, a more optimal block decomposition is suggested in Ref. [396], which we discuss in Section 10.2.4. It should be noted that the more optimal decomposition weakly breaks the reflection symmetry inherent to most global analog quantum simulation experiments, and thus can lead to small coupling between the symmetric and antisymmetric subspaces; if signals of interest are sensitive to this divide, such as preparation of GHZ states, then use caution and be sure to use a small enough δt , and a large enough block size.

With this approach – which we alternatively call ‘Lightcone’ simulation, or ‘ED-Haah’ (after the first author of Ref. [396]) – we are able to simulate systems of up to 28 qubits. Further, the Lightcone algorithm is easily adaptable to a matrix product state representation, which we find to be the highest fidelity simulation method for one-dimensional Rydberg atom arrays (Chapter 10).

B.3 RANDOM UNITARY CIRCUITS

Finally, here we describe random unitary circuit (RUC) simulation, as we used extensively throughout this thesis. In general random unitary circuits are not strictly defined: they could be a combination of randomly chosen single qubit gates followed by fixed two-qubit gates (as

for instance used in the Google supremacy paper [162]), or they could be composed of two-qubit Haar-randomized SU(4) gates, as are considered in tests of quantum volume [277]. Here we shall just describe the SU(4) implementation that we employ throughout this thesis.

The dynamics of the one-dimensional random unitary circuit (RUC) shown in are simulated using randomly sampled two-qubit SU(4) unitary gates from the Haar measure¹. Specifically, the time evolution of an N -qubit system starting from the initial state $|0\rangle^{\otimes N}$ can be described as follows:

$$|\psi(t)\rangle = \hat{\mathcal{U}}_t \hat{\mathcal{U}}_{t-1} \cdots \hat{\mathcal{U}}_2 \hat{\mathcal{U}}_1 |0\rangle^{\otimes N}, \quad (\text{B.8})$$

where $\hat{\mathcal{U}}_{\text{odd}} = \{\hat{\mathcal{U}}_1, \hat{\mathcal{U}}_3, \dots\}$ and $\hat{\mathcal{U}}_{\text{even}} = \{\hat{\mathcal{U}}_2, \hat{\mathcal{U}}_4, \dots\}$ are the odd- and even-time unitaries composed of local two-qubit unitaries as

$$\hat{\mathcal{U}}_{\text{odd}} = \prod_{i=1}^{N/2} \hat{U}_{2i-1,2i}, \quad \hat{\mathcal{U}}_{\text{even}} = \prod_{i=1}^{N/2} \hat{U}_{2i,2i+1} \quad (\text{B.9})$$

with open boundary condition, and t is the circuit depth and $\hat{U}_{\mu,\nu}$ is the randomly-sampled SU(4) gate acting on two qubits at site μ and ν . Note that at each circuit depth t , we randomly sample two-qubit random unitaries to generate a many-body unitary $\hat{\mathcal{U}}$, which leads to chaotic dynamics.

To emulate noisy quantum dynamics, we utilize a stochastic evolution method [114] incorporating local errors represented by the Pauli operators $\hat{S}^{x,y,z}$. These local errors are stochastically applied to individual qubits at a given single-qubit error rate per circuit layer. We repeat the noisy dynamics for thousands of trajectories at a fixed error rate, and further repeat over around 10 different circuit realizations. Through this repetition, we obtain a good approximation of the density matrix, $\hat{\rho}_{\text{RUC}}$.

¹To generate a randomized SU(D) gate the procedure [484] is:

1. Generate a $D \times D$ complex, Gaussian random matrix, i.e. one in which each element is given by $(x+iy)/\sqrt{2}$ where x and y are Gaussian random numbers with zero mean and unity standard deviation.
2. Perform QR decomposition on the matrix, yielding matrices Q and R .
3. Normalize $R' = \text{diag}(R)/|\text{diag}(R)|$, where diag picks out the diagonal elements of R .
4. The SU(D) gate is then $Q \times R'$.

Many-body temperature

In many of the studies we have conducted, it was important to define a notion of ‘temperature’ for many-body quantum states (for instance when discussing random state ensembles in Chapters 5 and 7, and for application of our benchmarking protocol in Chapters 6 and 8). Here we describe what we mean by temperature, and specifically describe observables of Rydberg atom arrays evolving at infinite temperature.

C.1 TEMPERATURE

As discussed extensively in Chapter 7, the developing study of *quantum thermalization* [305, 307, 308, 325–329] has sought to create a consistent formulation for both statistical and quantum mechanics. See Ref. [329] for the most pedagogical introduction to this topic. With this comes a notion of temperature for generic many-body quantum states; we will employ two different methods to practically define this temperature, one based on matching the energy density of a thermal state, and one based directly on the entropy of the quantum state.

As a disclaimer for both methods, it is very important to only study temperature in the irreducible Hilbert space of interest. As a concrete example, often in this thesis we considered states equilibrated at infinite effective temperature in the blockaded Hilbert space. When performing calculations to determine the temperature it is important to truncate the Hilbert space *only* to the blockaded manifold, because while a state may be at infinite temperature with respect to the blockade manifold it will be at very low temperature with respect to the total Hilbert space. The same care must be taken if there are symmetries in the system

(i.e. reflection, translation, etc.), which must be eliminated by writing the Hamiltonian in its irreducible representation (see Appendix B).

Energy density

Consider a pure quantum state, $|\psi_0\rangle$, which can be expressed in terms of the eigenvalues, E_n , and eigenstates, $|E_n\rangle$ of a given Hamiltonian, \hat{H} , as $|\psi_0\rangle = \sum_n c_n |E_n\rangle$ for coefficients c_n . The corresponding density matrix is $\hat{\rho} = |\psi_0\rangle\langle\psi_0| = \sum_{n,m} c_n c_m^* |E_n\rangle\langle E_m|$. We could alternatively consider a mixed state, $\hat{\rho}_\beta$, which is an incoherent (i.e. diagonal) mixture of the eigenstates, $\hat{\rho}_\beta = \sum_n p_n |E_n\rangle\langle E_n|$, for probabilities p_n with $\sum_n p_n = 1$.

To make connection to classical statistical mechanics, we then consider the mixture of eigenstates to describe a canonical ensemble at a temperature T (and inverse temperature $\beta = 1/T$), for which the probabilities are defined via Boltzmann factors $p_n = \exp(-\beta E_n)$ (we set Boltzmann's constant to 1 throughout). Thus, the thermal density matrix is written as

$$\hat{\rho}_\beta \propto \sum_n e^{-\beta E_n} |E_n\rangle\langle E_n| = e^{-\beta \hat{H}}, \quad (\text{C.1})$$

where the thermal expectation value of any operator, \hat{O} , can be written simply as

$$\langle \hat{O} \rangle = \frac{\text{Tr}[\hat{O} e^{-\beta \hat{H}}]}{\text{Tr}[e^{-\beta \hat{H}}]}. \quad (\text{C.2})$$

So far, we have only been making definitional statements. Now we make the explicit claim that the effective temperature of the pure quantum state $|\psi_0\rangle$ is the value of β for which the expectation value of the energy is the same for $|\psi_0\rangle$ and for $\hat{\rho}_\beta$. In other words, we equate

$$\langle \psi_0 | \hat{H} | \psi_0 \rangle = \frac{\text{Tr}[\hat{H} e^{-\beta \hat{H}}]}{\text{Tr}[e^{-\beta \hat{H}}]}. \quad (\text{C.3})$$

Thus, for any initial state, we can numerically test different values of β until Eq. (C.3) is fulfilled, at which point we associate β with the effective temperature of the state.

Then, if we for instance want to design a Hamiltonian such that the initial state is at infinite effective temperature, we simply pick a Hamiltonian parameter (e.g. detuning) to tune, until we find the point where $\beta = 0$. Note that such an infinite temperature point formally is one for which the overlap of the initial state with all energy eigenstates is equal, but in general for quantum many-body systems there will be some 'speckle' around this condition, even for (nearly) infinite effective temperature evolution.

Entropy

The method described above is a good conceptual way to define temperature, but it requires explicitly diagonalizing the Hamiltonian, and thus is limited to reasonably small system sizes. Here we describe a different, equivalent method of defining temperature directly from the entropy of a given quantum state, which can be modified to be performed with larger systems.

As a reminder, quantum thermalization deals with the situation of a bipartite quantum system, where one subsystem is the system of interest, while the other acts as an effectively thermal bath. While there is not an actual notion of thermal entropy for pure quantum states, we instead associate entanglement entropy between the two subsystems as analogous to thermal entropy. For a given density matrix, $\hat{\rho}$ we can write the entropy as

$$S = -\text{Tr}[\hat{\rho} \log(\hat{\rho})]. \quad (\text{C.4})$$

From classical thermodynamics we can then relate the entropy to the temperature via the simple differential relationship

$$\delta S = \beta \delta E \quad (\text{C.5})$$

where δE is a differential change in energy. To utilize this definition, we consider the diagonal density matrix (explored in greater detail in Chapter 7), which takes the form

$$\hat{\rho}_d = \sum_n |c_n|^2 |E_n\rangle \langle E_n|, \quad (\text{C.6})$$

where $|E_n\rangle$ are energy eigenstates, and $c_n = |\langle E_n | \psi_0 \rangle|^2$. If we normalize $\hat{\rho}_d$ in the irreducible representation of the Hilbert space (i.e. accounting for all symmetries, blockade, etc), then we may simply calculate the entropy as $S = -\text{Tr}[\hat{\rho}_d \log(\hat{\rho}_d)]$.

As before, if we wish to find an infinite temperature condition, we can tune a term of the Hamiltonian until we find the entropy is maximized with respect to the tuning parameter: at this point the slope of the entropy with respect to the system energy is 0, and so the effective temperature is infinite. This method will be consistent with the thermal density matrix approach above, but in general yields broader features. Also, it still explicitly requires us to diagonalize the Hamiltonian to calculate $\hat{\rho}_d$.

To get around this problem, we could instead obtain $\hat{\rho}_d$ directly as

$$\hat{\rho}_d = \lim_{\tau \rightarrow \infty} \int_0^\tau dt |\psi(t)\rangle \langle \psi(t)| \quad (\text{C.7})$$

where $|\psi(t)\rangle$ is the time-evolved initial state. One can perform this time-evolution for relatively large system sizes out to a very long time, then numerically construct the $\hat{\rho}_d$. After truncating to the appropriate irreducible representation, the same entropy-based analysis above applies.

C.2 INFINITE TEMPERATURE RYDBERG ATOM ARRAYS

When systems are equilibrated at infinite effective temperature, a number of their key characteristics may be simplified. First, we better define the blockaded Hilbert space dimension, and then describe calculating certain observables for a Rydberg atom array at infinite effective temperature.

The blockaded Hilbert space dimension

It is somewhat common knowledge that the dimension of the blockaded Hilbert space for N atoms with open boundary conditions is given by $\text{Fib}(N+2) \approx \varphi^N$, where $\varphi = \frac{1+\sqrt{5}}{2} \approx 1.6$ is the golden ratio. However, it is nice to see this derived. To do so, consider a chain of $N-1$ atoms. Let's label the last atom of this $N-1$ atom chain as x , and consider adding one more atom, which we'll call y , to its right. If $y=0$, then the Hilbert space dimension is the same as that of the entire $N-1$ atom chain. However, if $y=1$ then we force $x=0$, and so the Hilbert space dimension is that of the $N-2$ atom chain. Thus, the Hilbert space dimension of the N qubit system will be $\dim(\mathcal{H}_N^{\text{OBC}}) = \dim(\mathcal{H}_{N-1}^{\text{OBC}}) + \dim(\mathcal{H}_{N-2}^{\text{OBC}})$. While this already looks very Fibonacci-ish, we can formalize this by writing the recurrence relationship in matrix form as

$$\begin{bmatrix} \dim(\mathcal{H}_N^{\text{OBC}}) \\ \dim(\mathcal{H}_{N-1}^{\text{OBC}}) \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \dim(\mathcal{H}_{N-1}^{\text{OBC}}) \\ \dim(\mathcal{H}_{N-2}^{\text{OBC}}) \end{bmatrix}. \quad (\text{C.8})$$

The eigenvalues of this matrix are the golden ratio and its negative inverse, as expected.

Note that for periodic boundary conditions, we instead have $\dim(\mathcal{H}_N^{\text{PBC}}) = \dim(\mathcal{H}_{N-1}^{\text{OBC}}) + \dim(\mathcal{H}_{N-3}^{\text{OBC}})$ (note the change from $(N-2) \rightarrow (N-3)$). This is because if the y atom is a 1, it not only restricts the x atom, but also the 'first' atom in the chain.

Rydberg expectation values at infinite temperature

Operating at an infinite effective temperature allows us to make analytical predictions for many experimental observables. For instance, what is the steady state Rydberg occupation (expectation value of \hat{n}) for such a system? To simplify our lives, we consider periodic boundary conditions, which is equivalent to working in the bulk of an open boundary

system, though the same types of calculations we will perform could be done for the edges of open boundary systems.

To calculate $\langle \hat{n} \rangle$ at a given arbitrary site, we note it is the number of blockaded bitstrings with a ‘1’ on a given site, divided by the number of total bitstrings. If there is a ‘1’ on a given site, we know its neighbors must be ‘0,’ meaning there are only $\dim(\mathcal{H}_{N-3}^{\text{OBC}})$ bitstrings satisfying this requirement. Thus, the expectation value will be

$$\langle \hat{n} \rangle = \frac{\dim(\mathcal{H}_{N-3}^{\text{OBC}})}{\dim(\mathcal{H}_N^{\text{PBC}})} \quad (\text{C.9})$$

$$= \frac{\varphi^{N-3}}{\varphi^{N-3} + \varphi^{N-1}} \quad (\text{C.10})$$

$$= \frac{1}{\varphi^2 + 1} \quad (\text{C.11})$$

$$\approx 0.276. \quad (\text{C.12})$$

The same analysis can be carried out for higher-order operators, like the anti-ferromagnetic connected two-point correlator $(-1)^{|i-j|}(\langle \hat{n}_i \hat{n}_j \rangle - \langle \hat{n}_i \rangle \langle \hat{n}_j \rangle)$ we considered in Chapter 11, for which we would find the infinite temperature prediction as a function of distance, $d = |i - j|$, is $0.2 \left(\frac{\varphi-1}{\varphi} \right)^d$.

Note that all these claims can be easily checked numerically by explicitly constructing the blockaded basis.

Bibliography

- [1] A. Einstein, B. Podolsky, and N. Rosen, Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?, *Physical Review* **47**, 777 (1935).
- [2] J. S. Bell, On the Einstein Podolsky Rosen paradox, *Physics Physique Fizika* **1**, 195 (1964).
- [3] S. J. Freedman and J. F. Clauser, Experimental Test of Local Hidden-Variable Theories, *Physical Review Letters* **28**, 938 (1972).
- [4] A. Aspect, J. Dalibard, and G. Roger, Experimental Test of Bell's Inequalities Using Time-Varying Analyzers, *Physical Review Letters* **49**, 1804 (1982).
- [5] R. P. Feynman, Simulating physics with computers, *International Journal of Theoretical Physics* **21**, 467 (1982).
- [6] Y. Manin, *Computable and Uncomputable*, vol. 128 of *Sovetskoye Radio* (Moscow, 1980).
- [7] P. Benioff, The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines, *Journal of Statistical Physics* **22**, 563 (1980).
- [8] A. Y. Kitaev, Quantum measurements and the Abelian Stabilizer Problem, Preprint at <http://arxiv.org/abs/quant-ph/9511026> (1995).
- [9] E. Bernstein and U. Vazirani, Quantum Complexity Theory, *SIAM Journal on Computing* **26**, 1411 (1997).
- [10] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 1996), STOC '96, pp. 212–219, ISBN 978-0-89791-785-8.

- [11] P. W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994), pp. 124–134.
- [12] P. W. Shor, Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer, *SIAM Journal on Computing* **26**, 1484 (1997).
- [13] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Physical Review Letters* **103**, 150502 (2009).
- [14] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica, et al., Quantum algorithms: A survey of applications and end-to-end complexities, Preprint at <http://arxiv.org/abs/2310.03011> (2023).
- [15] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010), ISBN 978-0-511-97666-7.
- [16] W. G. Unruh, Maintaining coherence in quantum computers, *Physical Review A* **51**, 992 (1995).
- [17] R. Landauer, Is Quantum Mechanics Useful?, *Philosophical Transactions: Physical Sciences and Engineering* **353**, 367 (1995).
- [18] M. Born, Zur Quantenmechanik der Stoßvorgänge, *Zeitschrift für Physik* **37**, 863 (1926).
- [19] J. A. Wheeler and W. H. Zurek, eds., *Quantum Theory and Measurement* (Princeton University Press, 1983).
- [20] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, *Physical Review A* **52**, R2493 (1995).
- [21] S. J. Devitt, W. J. Munro, and K. Nemoto, Quantum error correction for beginners, *Reports on Progress in Physics* **76**, 076001 (2013).
- [22] J. Roffe, Quantum error correction: an introductory guide, *Contemporary Physics* **60**, 226 (2019).
- [23] A. Kitaev, Fault-tolerant quantum computation by anyons, *Annals of Physics* **303**, 2 (2003).
- [24] E. Knill, R. Laflamme, and W. H. Zurek, Resilient Quantum Computation, *Science* **279**, 342 (1998).
- [25] D. Aharonov and M. Ben-Or, Fault-Tolerant Quantum Computation with Constant Error Rate, *SIAM Journal on Computing* **38**, 1207 (2008).
- [26] P. Scholl, A. L. Shaw, R. B.-S. Tsai, R. Finkelstein, J. Choi, and M. Endres, Erasure conversion in a high-fidelity Rydberg quantum simulator, *Nature* **622**, 273 (2023).

- [27] A. L. Shaw, D. K. Mark, J. Choi, R. Finkelstein, P. Scholl, S. Choi, and M. Endres, Universal fluctuations and noise learning from Hilbert-space ergodicity, Preprint at <http://arxiv.org/abs/2403.11971> (2024).
- [28] A. L. Shaw, R. Finkelstein, R. B.-S. Tsai, P. Scholl, T. H. Yoon, J. Choi, and M. Endres, Multi-ensemble metrology by programming local rotations with atom movements, *Nature Physics* **20**, 195 (2024).
- [29] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale, et al., Preparing random states and benchmarking with many-body quantum chaos, *Nature* **613**, 468 (2023).
- [30] D. K. Mark, J. Choi, A. L. Shaw, M. Endres, and S. Choi, Benchmarking Quantum Simulators Using Ergodic Quantum Dynamics, *Physical Review Letters* **131**, 110601 (2023).
- [31] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, Benchmarking highly entangled states on a 60-atom analog quantum simulator, Preprint at <http://arxiv.org/abs/2308.07914>, in press at *Nature* (2023).
- [32] P. Scholl, A. L. Shaw, R. Finkelstein, R. B.-S. Tsai, J. Choi, and M. Endres, Erasure-cooling, control, and hyper-entanglement of motion in optical tweezers, Preprint at <http://arxiv.org/abs/2311.15580> (2023).
- [33] A. L. Shaw, P. Scholl, R. Finkelstein, I. S. Madjarov, B. Grinkemeyer, and M. Endres, Dark-State Enhanced Loading of an Optical Tweezer Array, *Physical Review Letters* **130**, 193402 (2023).
- [34] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, High-fidelity entanglement and detection of alkaline-earth Rydberg atoms, *Nature Physics* **16**, 857 (2020).
- [35] A. M. Kaufman and K.-K. Ni, Quantum science with optical tweezer arrays of ultracold atoms and molecules, *Nature Physics* **17**, 1324 (2021).
- [36] A. Browaeys and T. Lahaye, Many-body physics with individually controlled Rydberg atoms, *Nature Physics* **16**, 132 (2020).
- [37] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, et al., A quantum processor based on coherent transport of entangled atom arrays, *Nature* **604**, 451 (2022).
- [38] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, et al., Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, *Nature* **604**, 457 (2022).
- [39] R. Finkelstein, R. B.-S. Tsai, X. Sun, P. Scholl, S. Direkci, T. Gefen, J. Choi, A. L. Shaw, and M. Endres, Universal quantum operations and ancilla-based readout for tweezer clocks, Preprint at <http://arxiv.org/abs/2402.16220> (2024).

- [40] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, et al., Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [41] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, et al., Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [42] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, et al., Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature* **595**, 233 (2021).
- [43] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, et al., Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature* **595**, 227 (2021).
- [44] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys, Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms, *Science* **365**, 775 (2019).
- [45] P. Scholl, H. J. Williams, G. Bornet, F. Wallner, D. Barredo, L. Henriët, A. Signoles, C. Hainaut, T. Franz, S. Geier, et al., Microwave Engineering of Programmable XXZ Hamiltonians in Arrays of Rydberg Atoms, *PRX Quantum* **3**, 020303 (2022).
- [46] G. Bornet, G. Emperauger, C. Chen, B. Ye, M. Block, M. Bintz, J. A. Boyd, D. Barredo, T. Comparin, F. Mezzacapo, et al., Scalable spin squeezing in a dipolar Rydberg atom array, *Nature* **621**, 728 (2023).
- [47] K. Kim, F. Yang, K. Mølmer, and J. Ahn, Realization of an Extremely Anisotropic Heisenberg Magnet in Rydberg Atom Arrays, *Physical Review X* **14**, 011025 (2024).
- [48] G. Birkhoff, in *A history of scientific computing* (Association for Computing Machinery, New York, NY, USA, 1990), pp. 63–87, ISBN 978-0-201-50814-7.
- [49] A. Y. Kitaev, Quantum computations: algorithms and error correction, *Russian Mathematical Surveys* **52**, 1191 (1997).
- [50] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [51] A. Botea, A. Kishimoto, and R. Marinescu, On the Complexity of Quantum Circuit Compilation, *Proceedings of the International Symposium on Combinatorial Search* **9**, 138 (2018).
- [52] B. Tan and J. Cong, Optimality Study of Existing Quantum Computing Layout Synthesis Tools, *IEEE Transactions on Computers* **70**, 1363 (2021).
- [53] P. Gokhale, A. Javadi-Abhari, N. Earnest, Y. Shi, and F. T. Chong, Optimized Quantum Compilation for Near-Term Algorithms with OpenPulse, in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)* (2020), pp. 186–200.

- [54] J. Kalloor, M. Weiden, E. Younis, J. Kubiawicz, B. De Jong, and C. Iancu, Quantum Hardware Roofline: Evaluating the Impact of Gate Expressivity on Quantum Processor Design, Preprint at <http://arxiv.org/abs/2403.00132> (2024).
- [55] Z. Cai and S. C. Benjamin, Constructing Smaller Pauli Twirling Sets for Arbitrary Error Channels, *Scientific Reports* **9**, 11281 (2019).
- [56] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Mixed-state entanglement and quantum error correction, *Physical Review A* **54**, 3824 (1996).
- [57] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Symmetrized Characterization of Noisy Quantum Processes, *Science* **317**, 1893 (2007).
- [58] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Physical Review A* **77**, 012307 (2008).
- [59] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, *Nature* **607**, 667 (2022).
- [60] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, et al., Quantum Simulators: Architectures and Opportunities, *PRX Quantum* **2**, 17003 (2021).
- [61] F. Flamini, N. Spagnolo, and F. Sciarrino, Photonic quantum information processing: a review, *Reports on Progress in Physics* **82**, 016001 (2019).
- [62] M. J. Hartmann, Quantum simulation with interacting photons, *Journal of Optics* **18**, 104005 (2016).
- [63] A. A. Houck, H. E. Türeci, and J. Koch, On-chip quantum simulation with superconducting circuits, *Nature Physics* **8**, 292 (2012).
- [64] C. Monroe, W. Campbell, L.-M. Duan, Z.-X. Gong, A. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, G. Pagano, et al., Programmable quantum simulations of spin systems with trapped ions, *Reviews of Modern Physics* **93**, 25001 (2021).
- [65] C. Gross and I. Bloch, Quantum simulations with ultracold atoms in optical lattices, *Science* **357**, 995 (2017).
- [66] R. Trivedi, A. F. Rubio, and J. I. Cirac, Quantum advantage and stability to errors in analogue quantum simulators, Preprint at <http://arxiv.org/abs/2212.04924> (2023).
- [67] M. Grassl, T. Beth, and T. Pellizzari, Codes for the quantum erasure channel, *Physical Review A* **56**, 33 (1997).
- [68] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, *Nature Communications* **13**, 4657 (2022).

- [69] I. S. Madjarov, A. Cooper, A. L. Shaw, J. P. Covey, V. Schkolnik, T. H. Yoon, J. R. Williams, and M. Endres, An Atomic-Array Optical Clock with Single-Atom Readout, *Physical Review X* **9**, 41052 (2019).
- [70] J. S. Cotler, D. K. Mark, H.-Y. Huang, F. Hernández, J. Choi, A. L. Shaw, M. Endres, and S. Choi, Emergent Quantum State Designs from Individual Many-Body Wave Functions, *PRX Quantum* **4**, 10311 (2023).
- [71] D. K. Mark, F. Surace, A. Elben, A. L. Shaw, J. Choi, G. Refael, M. Endres, and S. Choi, A Maximum Entropy Principle in Deep Thermalization and in Hilbert-Space Ergodicity, Preprint at <http://arxiv.org/abs/2403.11970> (2024).
- [72] G. Vidal and R. F. Werner, Computable measure of entanglement, *Physical Review A* **65**, 032314 (2002).
- [73] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, Preprint at <http://arxiv.org/abs/2403.12021> (2024).
- [74] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, et al., High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [75] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, High-fidelity gates and mid-circuit erasure conversion in an atomic qubit, *Nature* **622**, 279 (2023).
- [76] I. S. Madjarov, Ph.D. thesis, California Institute of Technology (2021).
- [77] H. Levine, Ph.D. thesis, Harvard University (2021).
- [78] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, G. Pagano, et al., Programmable Quantum Simulations of Spin Systems with Trapped Ions, *arXiv:1912.07845* (2019).
- [79] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm, and S. Chu, Observation of a single-beam gradient force optical trap for dielectric particles, *Optics Letters* **11**, 288 (1986).
- [80] A. Ashkin, Optical trapping and manipulation of neutral particles using lasers, *Proceedings of the National Academy of Sciences* **94**, 4853 (1997).
- [81] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, Atom-by-atom assembly of defect-free one-dimensional cold atom arrays, *Science* **354**, 1024 (2016).
- [82] P. Huft, Y. Song, T. M. Graham, K. Jooya, S. Deshpande, C. Fang, M. Kats, and M. Saffman, Simple, passive design for large optical trap arrays for single atoms, *Physical Review A* **105**, 063111 (2022).

- [83] L. Pause, L. Sturm, M. Mittenbühler, S. Amann, T. Preuschoff, D. Schäffner, M. Schlosser, and G. Birkl, Supercharged two-dimensional tweezer array with more than 1000 atomic qubits, *Optica* **11**, 222 (2024).
- [84] M. A. Norcia, H. Kim, W. B. Cairncross, M. Stone, A. Ryou, M. Jaffe, M. O. Brown, K. Barnes, P. Battaglino, T. C. Bohdanowicz, et al., Iterative assembly of ^{171}Yb atom arrays in cavity-enhanced optical lattices, Preprint at <http://arxiv.org/abs/2401.16177> (2024).
- [85] J. P. Covey, I. S. Madjarov, A. Cooper, and M. Endres, 2000-Times Repeated Imaging of Strontium Atoms in Clock-Magic Tweezer Arrays, *Physical Review Letters* **122**, 173201 (2019).
- [86] K.-N. Schymik, S. Pancaldi, F. Nogrette, D. Barredo, J. Paris, A. Browaeys, and T. Lahaye, Single Atoms with 6000-Second Trapping Lifetimes in Optical-Tweezer Arrays at Cryogenic Temperatures, *Physical Review Applied* **16**, 034013 (2021).
- [87] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, Single-atom-resolved fluorescence imaging of an atomic Mott insulator, *Nature* **467**, 68 (2010).
- [88] K. D. Nelson, X. Li, and D. S. Weiss, Imaging single atoms in a three-dimensional array, *Nature Physics* **3**, 556 (2007).
- [89] D. Barredo, S. de Leseleuc, V. Lienhard, T. Lahaye, and A. Browaeys, An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays, *Science* **354**, 1021 (2016).
- [90] K.-N. Schymik, V. Lienhard, D. Barredo, P. Scholl, H. Williams, A. Browaeys, and T. Lahaye, Enhanced atom-by-atom assembly of arbitrary tweezer arrays, *Physical Review A* **102**, 63107 (2020).
- [91] D. Barredo, V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys, Synthetic three-dimensional atomic structures assembled atom by atom, *Nature* **561**, 79 (2018).
- [92] W. Lee, H. Kim, and J. Ahn, Three-dimensional rearrangement of single atoms using actively controlled optical microtraps, *Optics Express* **24**, 9816 (2016).
- [93] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Fast Quantum Gates for Neutral Atoms, *Physical Review Letters* **85**, 2208 (2000).
- [94] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller, Dipole Blockade and Quantum Information Processing in Mesoscopic Atomic Ensembles, *Physical Review Letters* **87**, 37901 (2001).
- [95] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, et al., Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms, *Physical Review Letters* **123**, 170503 (2019).

- [96] S. Stastny, H. P. Büchler, and N. Lang, Functional completeness of planar Rydberg blockade structures, *Physical Review B* **108**, 085138 (2023).
- [97] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, et al., Controlling quantum many-body dynamics in driven Rydberg atom arrays, *Science* **371**, 1355 (2021).
- [98] P. Patil, Quantum Monte Carlo simulations in the restricted Hilbert space of Rydberg atom arrays, Preprint at <http://arxiv.org/abs/2309.00482> (2023).
- [99] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, Alkaline-Earth Atoms in Optical Tweezers, *Physical Review X* **8**, 41055 (2018).
- [100] M. A. Norcia, A. W. Young, and A. M. Kaufman, Microscopic Control and Detection of Ultracold Strontium in Optical-Tweezer Arrays, *Physical Review X* **8**, 41054 (2018).
- [101] R. Tao, M. Ammenwerth, F. Gyger, I. Bloch, and J. Zeiher, High-fidelity detection of large-scale atom arrays in an optical lattice, Preprint at <http://arxiv.org/abs/2309.04717> (2023).
- [102] A. Urech, I. H. A. Knottnerus, R. J. C. Spreeuw, and F. Schreck, Narrow-line imaging of single strontium atoms in shallow optical tweezers, *Physical Review Research* **4**, 023245 (2022).
- [103] S. Saskin, J. T. Wilson, B. Grinkemeyer, and J. D. Thompson, Narrow-Line Cooling and Imaging of Ytterbium Atoms in an Optical Tweezer Array, *Physical Review Letters* **122**, 143002 (2019).
- [104] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman, Ytterbium Nuclear-Spin Qubits in an Optical Tweezer Array, *Physical Review X* **12**, 21027 (2022).
- [105] W. Huie, L. Li, N. Chen, X. Hu, Z. Jia, W. K. C. Sun, and J. P. Covey, Repetitive Readout and Real-Time Control of Nuclear Spin Qubits in ^{171}Yb Atoms, *PRX Quantum* **4**, 030337 (2023).
- [106] M. Norcia, W. Cairncross, K. Barnes, P. Battaglino, A. Brown, M. Brown, K. Cas-sella, C.-A. Chen, R. Coxe, D. Crow, et al., Midcircuit Qubit Measurement and Rearrangement in a ^{171}Yb Atomic Array, *Physical Review X* **13**, 041034 (2023).
- [107] S. Stellmer, F. Schreck, and T. C. Killian, Degenerate quantum gases of strontium, *Annual Review of Cold Atoms and Molecules* **1**, 1 (2014).
- [108] R. Santra, K. V. Christ, and C. H. Greene, Properties of metastable alkaline-earth-metal atoms calculated using an accurate effective core potential, *Physical Review A* **69**, 42510 (2004).

- [109] A. Taichenachev, V. Yudin, C. Oates, C. Hoyt, Z. Barber, and L. Hollberg, Magnetic Field-Induced Spectroscopy of Forbidden Optical Transitions with Application to Lattice-Based Optical Atomic Clocks, *Physical Review Letters* **96**, 83001 (2006).
- [110] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, Resolving the gravitational redshift across a millimetre-scale atomic sample, *Nature* **602**, 420 (2022).
- [111] X. Zheng, J. Dolde, V. Lochab, B. N. Merriman, H. Li, and S. Kolkowitz, Differential clock comparisons with a multiplexed optical lattice clock, *Nature* **602**, 425 (2022).
- [112] M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman, Seconds-scale coherence on an optical clock transition in a tweezer array, *Science* **366**, 93 (2019).
- [113] A. P. Burgers, S. Ma, S. Saskin, J. Wilson, M. A. Alarcón, C. H. Greene, and J. D. Thompson, Controlling Rydberg Excitations Using Ion-Core Transitions in Alkaline-Earth Atom-Tweezer Arrays, *PRX Quantum* **3**, 020326 (2022).
- [114] K. Mølmer, Y. Castin, and J. Dalibard, Monte Carlo wave-function method in quantum optics, *Journal of the Optical Society of America B* **10**, 524 (1993).
- [115] S. de Léséleuc, D. Barredo, V. Lienhard, A. Browaeys, and T. Lahaye, Analysis of imperfections in the coherent optical excitation of single atoms to Rydberg states, *Physical Review A* **97**, 53803 (2018).
- [116] E. D. Black, An introduction to Pound–Drever–Hall laser frequency stabilization, *American Journal of Physics* **69**, 79 (2001).
- [117] A. M. Kale, Senior Thesis, California Institute of Technology (2020).
- [118] W. Nagourney, *Quantum Electronics for Atomic Physics and Telecommunication* (OUP Oxford, 2014), ISBN 0-19-164338-6.
- [119] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman, Long-lived Bell states in an array of optical clock qubits, *Nature Physics* **18**, 1067 (2022).
- [120] M. Chafer, J. Osório, A. Dhaybi, F. Ravetta, F. Amrani, B. Debord, C. Cailteau-Fischbach, F. Gérôme, G. Ancellet, and F. Benabid, Ultra-compact 266-289 nm pair source for DIAL LIDAR based on hollow-core photonic crystal fiber, in *2021 Conference on Lasers and Electro-Optics (CLEO)* (2021), pp. 1–3, iSSN: 2160-8989.
- [121] Y. Colombe, D. H. Slichter, A. C. Wilson, D. Leibfried, and D. J. Wineland, Single-mode optical fiber for high-power, low-loss UV transmission, *Optics Express* **22**, 19783 (2014).
- [122] C. D. Marciniak, H. B. Ball, A. T.-H. Hung, and M. J. Biercuk, Towards fully commercial, UV-compatible fiber patch cords, *Optics Express* **25**, 15643 (2017).

- [123] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, Optical Dipole Traps for Neutral Atoms, *Adv. Atom. Mol. Opt. Phys.* **42**, 95 (2000).
- [124] D. J. Wineland and W. M. Itano, Laser cooling of atoms, *Physical Review A* **20**, 1521 (1979).
- [125] W. D. Phillips, Nobel Lecture: Laser cooling and trapping of neutral atoms, *Reviews of Modern Physics* **70**, 721 (1998).
- [126] C. N. Cohen-Tannoudji, Nobel Lecture: Manipulating atoms with photons, *Reviews of Modern Physics* **70**, 707 (1998).
- [127] S. Chu, Nobel Lecture: The manipulation of neutral particles, *Reviews of Modern Physics* **70**, 685 (1998).
- [128] H. J. Metcalf and P. Van Der Straten, *Laser Cooling and Trapping*, Graduate Texts in Contemporary Physics (Springer, New York, NY, 1999), ISBN 978-0-387-98728-6 978-1-4612-1470-0.
- [129] D. A. Steck, *Quantum and Atom Optics* (available online at <http://steck.us/teaching>, 2007).
- [130] A. M. Kaufman, B. J. Lester, and C. A. Regal, Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State, *Physical Review X* **2**, 41014 (2012).
- [131] K. Gillen-Christandl, G. D. Gillen, M. J. Piotrowicz, and M. Saffman, Comparison of Gaussian and super Gaussian laser beams for addressing atomic qubits, *Applied Physics B* **122**, 131 (2016).
- [132] K. N. Blodgett, D. Peana, S. S. Phatak, L. M. Terry, M. P. Montes, and J. D. Hood, Imaging a Li Atom in an Optical Tweezer 2000 Times with Lambda-Enhanced Gray Molasses, *Physical Review Letters* **131**, 083001 (2023).
- [133] S. Dörscher, R. Schwarz, A. Al-Masoudi, S. Falke, U. Sterr, and C. Lisdat, Lattice-induced photon scattering in an optical lattice clock, *Physical Review A* **97**, 63419 (2018).
- [134] J. Winklmann, D. Tsevas, and M. Schulz, Realistic Neutral Atom Image Simulation, in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2023), vol. 01, pp. 1349–1359.
- [135] A. Impertro, J. F. Wienand, S. Häfele, H. von Raven, S. Hubele, T. Klostermann, C. R. Cabrera, I. Bloch, and M. Aidelsburger, An unsupervised deep learning algorithm for single-site reconstruction in quantum gas microscopes, *Communications Physics* **6**, 1 (2023).
- [136] A. Cao, W. J. Eckner, T. L. Yelin, A. W. Young, S. Jandura, L. Yan, K. Kim, G. Pupillo, J. Ye, N. D. Oppong, et al., Multi-qubit gates and 'Schrodinger cat' states in an optical clock, Preprint at <http://arxiv.org/abs/2402.16289> (2024).

- [137] C. Nill, K. Brandner, B. Olmos, F. Carollo, and I. Lesanovsky, Many-Body Radiative Decay in Strongly Interacting Rydberg Ensembles, *Physical Review Letters* **129**, 243202 (2022).
- [138] A. Bergschneider, V. M. Klinkhamer, J. H. Becher, R. Klemt, G. Zürn, P. M. Preiss, and S. Jochim, Spin-resolved single-atom imaging of multiLi/mi mprescripts/ none/ mn6/mn /mmultiscripts /math in free space, *Physical Review A* **97**, 63613 (2018).
- [139] E. Deist, Y.-H. Lu, J. Ho, M. K. Pasha, J. Zeiher, Z. Yan, and D. M. Stamper-Kurn, Mid-Circuit Cavity Measurement in a Neutral Atom Array, *Physical Review Letters* **129**, 203602 (2022).
- [140] K. Singh, S. Anand, A. Pocklington, J. T. Kemp, and H. Bernien, Dual-Element, Two-Dimensional Atom Array with Continuous-Mode Operation, *Physical Review X* **12**, 11040 (2022).
- [141] F. Gyger, M. Ammenwerth, R. Tao, H. Timme, S. Snigirev, I. Bloch, and J. Zeiher, Continuous operation of large-scale atom arrays in optical lattices, Preprint at <http://arxiv.org/abs/2402.04994> (2024).
- [142] J. Samland, S. Bennetts, C.-C. Chen, R. G. Escudero, F. Schreck, and B. Pasquiou, Optical pumping of $5s4d$ $1D_2$ strontium atoms for laser cooling and imaging, Preprint at <http://arxiv.org/abs/2311.02941> (2024).
- [143] C. R. Clark, H. N. Tinkey, B. C. Sawyer, A. M. Meier, K. A. Burkhardt, C. M. Seck, C. M. Shappert, N. D. Guise, C. E. Volin, S. D. Fallek, et al., High-Fidelity Bell-State Preparation with 40Ca^+ Optical Qubits, *Physical Review Letters* **127**, 130505 (2021).
- [144] V. Negîrneac, H. Ali, N. Muthusubramanian, F. Battistel, R. Sagastizabal, M. Moreira, J. Marques, W. Vlothuizen, M. Beekman, C. Zachariadis, et al., High-Fidelity Controlled- Z Gate with Maximal Intermediate Leakage Operating at the Speed Limit in a Superconducting Quantum Processor, *Physical Review Letters* **126**, 220502 (2021).
- [145] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, *Applied Physics Reviews* **6**, 21314 (2019).
- [146] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting Qubits: Current State of Play, *Annual Review of Condensed Matter Physics* **11**, 369 (2020).
- [147] V. Giovannetti, S. Lloyd, and L. MacCone, Quantum metrology, *Physical Review Letters* **96**, 10401 (2006).
- [148] S. Pirandola, J. Eisert, C. Weedbrook, A. Furusawa, and S. L. Braunstein, Advances in quantum teleportation, *Nature Photonics* **9**, 641 (2015).

- [149] A. Gaëtan, Y. Miroshnychenko, T. Wilk, A. Chotia, M. Viteau, D. Comparat, P. Pillet, A. Browaeys, and P. Grangier, Observation of collective excitation of two individual atoms in the Rydberg blockade regime, *Nature Physics* **5**, 115 (2009).
- [150] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman, Demonstration of a Neutral Atom Controlled-NOT Quantum Gate, *Physical Review Letters* **104**, 10503 (2010).
- [151] D. Hayes, D. Stack, B. Bjork, A. Potter, C. Baldwin, and R. Stutz, Eliminating Leakage Errors in Hyperfine Qubits, *Physical Review Letters* **124**, 170501 (2020).
- [152] R. Stricker, D. Vodola, A. Erhard, L. Postler, M. Meth, M. Ringbauer, P. Schindler, T. Monz, M. Müller, and R. Blatt, Experimental deterministic correction of qubit loss, *Nature* **585**, 207 (2020).
- [153] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. J. Satzinger, C. Quintana, P. V. Klimov, D. Sank, C. Gidney, A. G. Fowler, et al., Removing leakage-induced correlated errors in superconducting quantum error correction, *Nature Communications* **12**, 1761 (2021).
- [154] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, High-Fidelity Control and Entanglement of Rydberg-Atom Qubits, *Physical Review Letters* **121**, 123603 (2018).
- [155] R. Löw, H. Weimer, J. Nipper, J. B. Balewski, B. Butscher, H. P. Büchler, and T. Pfau, An experimental and theoretical guide to strongly interacting Rydberg gases, *Journal of Physics B: Atomic, Molecular and Optical Physics* **45**, 113001 (2012).
- [156] P. Fendley, K. Sengupta, and S. Sachdev, Competing density-wave orders in a one-dimensional hard-boson model, *Physical Review B* **69**, 75106 (2004).
- [157] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, et al., Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).
- [158] R. Samajdar, S. Choi, H. Pichler, M. D. Lukin, and S. Sachdev, Numerical study of the chiral Z₃ quantum phase transition in one spatial dimension, *Physical Review A* **98**, 023614 (2018).
- [159] V. Eisler, Z. Rácz, and F. van Wijland, Magnetization distribution in the transverse Ising chain with energy flux, *Physical Review E* **67**, 056129 (2003).
- [160] A. Lamacraft and P. Fendley, Order Parameter Statistics in the Critical Quantum Ising Chain, *Physical Review Letters* **100**, 165706 (2008).
- [161] A. F. Albuquerque, F. Alet, C. Sire, and S. Capponi, Quantum critical scaling of fidelity susceptibility, *Physical Review B* **81**, 064418 (2010).

- [162] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, et al., Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [163] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, et al., Quantum optimization of maximum independent set using Rydberg atom arrays, *Science* **376**, 1209 (2022).
- [164] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Reviews of Modern Physics* **90**, 35005 (2018).
- [165] P. Niroula, J. Dolde, X. Zheng, J. Bringewatt, A. Ehrenberg, K. C. Cox, J. Thompson, M. J. Gullans, S. Kolkowitz, and A. V. Gorshkov, Quantum Sensing with Erasure Qubits, Preprint at <http://arxiv.org/abs/2310.01512> (2023).
- [166] H. J. Carmichael, Quantum trajectory theory for cascaded open systems, *Physical Review Letters* **70**, 2273 (1993).
- [167] B. Skinner, J. Ruhman, and A. Nahum, Measurement-Induced Phase Transitions in the Dynamics of Entanglement, *Physical Review X* **9**, 031009 (2019).
- [168] Y. Li, X. Chen, and M. P. A. Fisher, Quantum Zeno effect and the many-body entanglement transition, *Physical Review B* **98**, 205136 (2018).
- [169] S. J. Garratt, Z. Weinstein, and E. Altman, Measurements Conspire Nonlocally to Restructure Critical Quantum States, *Physical Review X* **13**, 021026 (2023).
- [170] C. Lee, S. C. Webster, J. Mosca Toba, O. Corfield, G. Porter, and R. C. Thompson, Measurement-based ground-state cooling of a trapped-ion oscillator, *Physical Review A* **107**, 033107 (2023).
- [171] X. Zhang, K. Beloy, Y. Hassan, W. McGrew, C.-C. Chen, J. Siegel, T. Grogan, and A. Ludlow, Subrecoil Clock-Transition Laser Cooling Enabling Shallow Optical Lattice Clocks, *Physical Review Letters* **129**, 113202 (2022).
- [172] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, Single-spin addressing in an atomic Mott insulator, *Nature* **471**, 319 (2011).
- [173] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen, Near-deterministic preparation of a single atom in an optical microtrap, *Nature Physics* **6**, 951 (2010).
- [174] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, Rapid Production of Uniformly Filled Arrays of Neutral Atoms, *Physical Review Letters* **115**, 73003 (2015).
- [175] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu, and C. A. Regal, Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly, *Physical Review X* **9**, 11057 (2019).

- [176] M. M. Aliyu, L. Zhao, X. Q. Quek, K. C. Yellapragada, and H. Loh, $\$D_1$ magic wavelength tweezers for scaling atom arrays, *Physical Review Research* **3**, 43059 (2021).
- [177] J. Ang'ong'a, C. Huang, J. P. Covey, and B. Gadway, Gray molasses cooling of 39K atoms in optical tweezers, *Physical Review Research* **4**, 13240 (2022).
- [178] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, Sub-poissonian loading of single atoms in a microscopic dipole trap, *Nature* **411**, 1024 (2001).
- [179] Y. H. Fung and M. F. Andersen, Efficient collisional blockade loading of a single atom into a tight microtrap, *New Journal of Physics* **17**, 73011 (2015).
- [180] L. Anderegg, L. W. Cheuk, Y. Bao, S. Burchesky, W. Ketterle, K.-K. Ni, and J. M. Doyle, An optical tweezer array of ultracold molecules, *Science* **365**, 1156 (2019).
- [181] C. M. Holland, Y. Lu, and L. W. Cheuk, On-Demand Entanglement of Molecules in a Reconfigurable Optical Tweezer Array, Preprint at <http://arxiv.org/abs/2210.06309> (2022).
- [182] C. W. Bauschlicher, S. R. Langhoff, and H. Partridge, The radiative lifetime of the 1 D 2 state of Ca and Sr: a core-valence treatment, *Journal of Physics B: Atomic and Molecular Physics* **18**, 1523 (1985).
- [183] T. Zelevinsky, M. M. Boyd, A. D. Ludlow, T. Ido, J. Ye, R. Ciuryło, P. Naidon, and P. S. Julienne, Narrow Line Photoassociation in an Optical Lattice, *Physical Review Letters* **96**, 203201 (2006).
- [184] G. Reinaudi, C. B. Osborn, M. McDonald, S. Kotochigova, and T. Zelevinsky, Optical Production of Stable Ultracold Sr₂ Molecules, *Physical Review Letters* **109**, 115303 (2012).
- [185] M. A. Norcia, Coupling atoms to cavities with narrow linewidth optical transitions: Applications to frequency metrology, arXiv:1908.11442 (2019).
- [186] A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman, Half-minute-scale atomic coherence and high relative stability in a tweezer clock, *Nature* **588**, 408 (2020).
- [187] N. Chen, L. Li, W. Huie, M. Zhao, I. Vetter, C. H. Greene, and J. P. Covey, Analyzing the Rydberg-based optical-metastable-ground architecture for 171 Yb nuclear spins, *Physical Review A* **105**, 052438 (2022).
- [188] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, Energy distribution and cooling of a single atom in an optical tweezer, *Physical Review A* **78**, 33425 (2008).
- [189] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais, and A. Browaeys, Free-Space Lossless State Detection of a Single Trapped Atom, *Physical Review Letters* **106**, 133003 (2011).

- [190] I. Cong, H. Levine, A. Keesling, D. Bluvstein, S.-T. Wang, and M. D. Lukin, Hardware-Efficient, Fault-Tolerant Quantum Computation with Rydberg Atoms, *Physical Review X* **12**, 21049 (2022).
- [191] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, Single-Atom Trapping in Holographic 2D Arrays of Microtraps with Arbitrary Geometries, *Physical Review X* **4**, 21034 (2014).
- [192] H. Kim, W. Lee, H.-g. Lee, H. Jo, Y. Song, and J. Ahn, In situ single-atom array synthesis using dynamic holographic optical tweezers, *Nature Communications* **7**, 13317 (2016).
- [193] M. A. Norcia, J. R. K. Cline, J. A. Muniz, J. M. Robinson, R. B. Hutson, A. Goban, G. E. Marti, J. Ye, and J. K. Thompson, Frequency Measurements of Superradiance from the Strontium Clock Transition, *Physical Review X* **8**, 21036 (2018).
- [194] Z. Z. Yan, B. M. Spar, M. L. Prichard, S. Chi, H.-T. Wei, E. Ibarra-García-Padilla, K. R. Hazzard, and W. S. Bakr, Two-Dimensional Programmable Tweezer Arrays of Fermions, *Physical Review Letters* **129**, 123201 (2022).
- [195] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, Universal Gate Operations on Nuclear Spin Qubits in an Optical Tweezer Array of ^{171}Yb Atoms, *Physical Review X* **12**, 021028 (2022).
- [196] T. M. Graham, L. Phuttitarn, R. Chinnarasu, Y. Song, C. Poole, K. Jooya, J. Scott, A. Scott, P. Eichler, and M. Saffman, Mid-circuit measurements on a single species neutral alkali atom quantum processor, Preprint at <http://arxiv.org/abs/2303.10051> (2023).
- [197] J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman, Midcircuit Operations Using the omg Architecture in Neutral Atom Arrays, *Physical Review X* **13**, 041035 (2023).
- [198] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. Schmidt, Optical atomic clocks, *Reviews of Modern Physics* **87**, 637 (2015).
- [199] E. Kessler, P. Kómár, M. Bishof, L. Jiang, A. Sørensen, J. Ye, and M. Lukin, Heisenberg-Limited Atom Clocks Based on Entangled Qubits, *Physical Review Letters* **112**, 190403 (2014).
- [200] T. Rosenband and D. R. Leibbrandt, Exponential scaling of clock stability with atom number (2013).
- [201] J. Borregaard and A. S. Sørensen, Efficient Atomic Clocks Operated with Several Atomic Ensembles, *Physical Review Letters* **111**, 90802 (2013).
- [202] B.-C. Ren and F.-G. Deng, Hyper-parallel photonic quantum computation with coupled quantum dots, *Scientific Reports* **4**, 4623 (2014).

- [203] F.-G. Deng, B.-C. Ren, and X.-H. Li, Quantum hyperentanglement and its applications in quantum information processing, *Science Bulletin* **62**, 46 (2017).
- [204] T. Đorđević, P. Samutpraphoot, P. L. Ocola, H. Bernien, B. Grinkemeyer, I. Dimitrova, V. Vuletić, and M. D. Lukin, Entanglement transport and a nanophotonic interface for atoms in optical tweezers, *Science* **373**, 1511 (2021).
- [205] A. Lengwenus, J. Kruse, M. Schlosser, S. Tichelmann, and G. Birkl, Coherent Transport of Atomic Quantum States in a Scalable Shift Register, *Physical Review Letters* **105**, 170502 (2010).
- [206] J. Beugnon, C. Tuchendler, H. Marion, A. Gaëtan, Y. Miroshnychenko, Y. R. P. Sortais, A. M. Lance, M. P. A. Jones, G. Messin, A. Browaeys, et al., Two-dimensional transport and transfer of a single atomic qubit in optical tweezers, *Nature Physics* **3**, 696 (2007).
- [207] T. Schaetz, M. D. Barrett, D. Leibfried, J. Chiaverini, J. Britton, W. M. Itano, J. D. Jost, C. Langer, and D. J. Wineland, Quantum Dense Coding with Atomic Qubits, *Physical Review Letters* **93**, 040505 (2004).
- [208] M. R. Lam, N. Peter, T. Groh, W. Alt, C. Robens, D. Meschede, A. Negretti, S. Montangero, T. Calarco, and A. Alberti, Demonstration of Quantum Brachistochrones between Distant States of an Atom, *Physical Review X* **11**, 011035 (2021).
- [209] W. Gong and S. Choi, Robust and Parallel Control of Many Qubits, Preprint at <http://arxiv.org/abs/2312.08426> (2023).
- [210] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss, Single-qubit gates based on targeted phase shifts in a 3D neutral atom array, *Science* **352**, 1562 (2016).
- [211] G. Wang, W. Xu, C. Li, V. Vuletić, and P. Cappellaro, Individual-atom control in array through phase modulation, Preprint at <http://arxiv.org/abs/2310.19741> (2023).
- [212] G. Wendin, Quantum information processing with superconducting circuits: a review, *Reports on Progress in Physics* **80**, 106001 (2017).
- [213] S. Slussarenko and G. J. Pryde, Photonic quantum information processing: A concise review, *Applied Physics Reviews* **6**, 041303 (2019).
- [214] P. K. Shandilya, S. Flågan, N. C. Carvalho, E. Zohari, V. K. Kavatamane, J. E. Losby, and P. E. Barclay, Diamond Integrated Quantum Photonics: A Review, *Journal of Lightwave Technology* **40**, 7538 (2022).
- [215] M. Kang, W. C. Campbell, and K. R. Brown, Quantum Error Correction with Metastable States of Trapped Ions Using Erasure Conversion, *PRX Quantum* **4**, 020358 (2023).
- [216] M. Saffman, Quantum computing with atomic qubits and Rydberg interactions: progress and challenges, *Journal of Physics B: Atomic, Molecular and Optical Physics* **49**, 202001 (2016).

- [217] M. O. Brown, S. R. Muleady, W. J. Dworschack, R. J. Lewis-Swan, A. M. Rey, O. Romero-Isart, and C. A. Regal, Time-of-flight quantum tomography of an atom in an optical tweezer, *Nature Physics* **19**, 569 (2023).
- [218] T. Hartke, B. Oreg, N. Jia, and M. Zwierlein, Quantum register of fermion pairs, *Nature* **601**, 537 (2022).
- [219] P. L. Knight, E. A. Hinds, M. B. Plenio, D. J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenković, et al., Quantum information processing with trapped ions, *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **361**, 1349 (2003).
- [220] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. P. Home, Encoding a qubit in a trapped-ion mechanical oscillator, *Nature* **566**, 513 (2019).
- [221] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, *Nature Communications* **8**, 94 (2017).
- [222] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, et al., Quantum error correction of a qubit encoded in grid states of an oscillator, *Nature* **584**, 368 (2020).
- [223] K. Singh, C. E. Bradley, S. Anand, V. Ramesh, R. White, and H. Bernien, Mid-circuit correction of correlated phase errors using an array of spectator qubits, *Science* **380**, 1265 (2023).
- [224] D. Gottesman, A. Kitaev, and J. Preskill, Encoding a qubit in an oscillator, *Physical Review A* **64**, 012310 (2001).
- [225] S. P. Jordan, K. S. M. Lee, and J. Preskill, Quantum computation of scattering in scalar quantum field theories, *Quantum Inf. Comput.* **14**, 1014 (2014).
- [226] A. Macridin, A. C. Y. Li, S. Mrenna, and P. Spentzouris, Bosonic field digitization for quantum computers, *Physical Review A* **105**, 052405 (2022).
- [227] B. M. Terhal, Quantum error correction for quantum memories, *Reviews of Modern Physics* **87**, 307 (2015).
- [228] J. Y. Lee, W. Ji, Z. Bi, and M. P. A. Fisher, Decoding Measurement-Prepared Quantum Phases and Transitions: from Ising model to gauge theory, and beyond, Preprint at <http://arxiv.org/abs/2208.11699> (2022).
- [229] N. Tantivasadakarn, R. Thorngren, A. Vishwanath, and R. Verresen, Long-range entanglement from measuring symmetry-protected topological phases, Preprint at <http://arxiv.org/abs/2112.01519> (2022).
- [230] X.-F. Shi, Hyperentanglement of divalent neutral atoms by Rydberg blockade, *Physical Review A* **104**, 042422 (2021).

- [231] X.-M. Hu, C.-X. Huang, Y.-B. Sheng, L. Zhou, B.-H. Liu, Y. Guo, C. Zhang, W.-B. Xing, Y.-F. Huang, C.-F. Li, et al., Long-Distance Entanglement Purification for Quantum Communication, *Physical Review Letters* **126**, 010503 (2021).
- [232] X.-M. Hu, Y. Guo, B.-H. Liu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, Beating the channel capacity limit for superdense coding with entangled ququarts, *Science Advances* **4**, eaat9304 (2018).
- [233] T. M. Graham, H. J. Bernstein, T.-C. Wei, M. Junge, and P. G. Kwiat, Superdense teleportation using hyperentangled photons, *Nature Communications* **6**, 7185 (2015).
- [234] H. Pichler, G. Zhu, A. Seif, P. Zoller, and M. Hafezi, Measurement Protocol for the Entanglement Spectrum of Cold Atoms, *Physical Review X* **6**, 041033 (2016).
- [235] G. M. Tino, Testing gravity with cold atom interferometry: results and prospects, *Quantum Science and Technology* **6**, 024014 (2021).
- [236] J. D. Jost, J. P. Home, J. M. Amini, D. Hanneke, R. Ozeri, C. Langer, J. J. Bollinger, D. Leibfried, and D. J. Wineland, Entangled mechanical oscillators, *Nature* **459**, 683 (2009).
- [237] S. Jandura and G. Pupillo, Time-Optimal Two- and Three-Qubit Gates for Rydberg Atoms, *Quantum* **6**, 712 (2022).
- [238] A. Pagano, S. Weber, D. Jaschke, T. Pfau, F. Meinert, S. Montangero, and H. P. Büchler, Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms, *Physical Review Research* **4**, 033019 (2022).
- [239] M. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, Search for new physics with atoms and molecules, *Reviews of Modern Physics* **90**, 25008 (2018).
- [240] V. Andreev, D. G. Ang, D. DeMille, J. M. Doyle, G. Gabrielse, J. Haefner, N. R. Hutzler, Z. Lasner, C. Meisenhelder, B. R. O’Leary, et al., Improved limit on the electric dipole moment of the electron, *Nature* **562**, 355 (2018).
- [241] T. S. Roussy, L. Caldwell, T. Wright, W. B. Cairncross, Y. Shagam, K. B. Ng, N. Schlossberger, S. Y. Park, A. Wang, J. Ye, et al., An improved bound on the electron’s electric dipole moment, *Science* **381**, 46 (2023).
- [242] C. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Reviews of Modern Physics* **89**, 35002 (2017).
- [243] E. Pedrozo-Peñafiel, S. Colombo, C. Shu, A. F. Adiyatullin, Z. Li, E. Mendez, B. Braverman, A. Kawasaki, D. Akamatsu, Y. Xiao, et al., Entanglement on an optical atomic-clock transition, *Nature* 2020 588:7838 **588**, 414 (2020).
- [244] K. Macieszczak, M. Fraas, and R. Demkowicz-Dobrzański, Bayesian quantum frequency estimation in presence of collective dephasing, *New Journal of Physics* **16**, 113002 (2014).

- [245] R. Kaubruegger, D. V. Vasilyev, M. Schulte, K. Hammerer, and P. Zoller, Quantum Variational Optimization of Ramsey Interferometry and Atomic Clocks, *Physical Review X* **11**, 41045 (2021).
- [246] S. F. Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac, Improvement of Frequency Standards with Quantum Entanglement, *Physical Review Letters* **79**, 3865 (1997).
- [247] M. Schulte, C. Lisdat, P. O. Schmidt, U. Sterr, and K. Hammerer, Prospects and challenges for squeezing-enhanced optical atomic clocks, *Nature Communications* 2020 11:1 **11**, 1 (2020).
- [248] S. Colombo, E. Pedrozo-Peñañiel, A. F. Adiyatullin, Z. Li, E. Mendez, C. Shu, and V. Vuletić, Time-reversal-based quantum metrology with many-body entangled states, *Nature Physics* **18**, 925 (2022).
- [249] V. Bužek, R. Derka, and S. Massar, Optimal Quantum Clocks, *Physical Review Letters* **82**, 2207 (1999).
- [250] W. Li, S. Wu, A. Smerzi, and L. Pezzè, Improved absolute clock stability by the joint interrogation of two atomic ensembles, *Physical Review A* **105**, 53116 (2022).
- [251] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland, Quantum projection noise: Population fluctuations in two-level systems, *Physical Review A* **47**, 3554 (1993).
- [252] W. Bowden, A. Vianello, I. R. Hill, M. Schioppo, and R. Hobson, Improving the Q Factor of an Optical Atomic Clock Using Quantum Nondemolition Measurement, *Physical Review X* **10**, 41052 (2020).
- [253] R. Kohlhaas, A. Bertoldi, E. Cantin, A. Aspect, A. Landragin, and P. Bouyer, Phase Locking a Clock Oscillator to a Coherent Atomic Ensemble, *Physical Review X* **5**, 021011 (2015).
- [254] A. Strathearn, R. F. Offer, A. P. Hilton, E. Klantsataya, A. N. Luiten, R. P. Anderson, B. M. Sparkes, and T. M. Stace, Wave-front curvature in optical atomic beam clocks, *Physical Review A* **108**, 013105 (2023).
- [255] C. D. Marciniak, T. Feldker, I. Pogorelov, R. Kaubruegger, D. V. Vasilyev, R. van Bijnen, P. Schindler, P. Zoller, R. Blatt, and T. Monz, Optimal metrology with programmable quantum sensors, *Nature* **603**, 604 (2022).
- [256] K. Barnes, P. Battaglino, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen, et al., Assembly and coherent control of a register of nuclear spin qubits, *Nature Communications* **13**, 2779 (2022).
- [257] A. Kubica, A. Haim, Y. Vaknin, H. Levine, F. Brandão, and A. Retzker, Erasure Qubits: Overcoming the $\{T\}_1$ Limit in Superconducting Circuits, *Physical Review X* **13**, 041022 (2023).

- [258] J. D. Teoh, P. Winkel, H. K. Babla, B. J. Chapman, J. Claes, S. J. de Graaf, J. W. O. Garmon, W. D. Kalfus, Y. Lu, A. Maiti, et al., Dual-rail encoding with superconducting cavities, *Proceedings of the National Academy of Sciences* **120**, e2221736120 (2023).
- [259] H. Levine, A. Haim, J. S. C. Hung, N. Alidoust, M. Kalae, L. DeLorenzo, E. A. Wollack, P. A. Arriola, A. Khalajhedayati, R. Sanil, et al., Demonstrating a long-coherence dual-rail erasure qubit using tunable transmons, Preprint at <http://arxiv.org/abs/2307.08737> (2023).
- [260] K. S. Chou, T. Shemma, H. McCarrick, T.-C. Chien, J. D. Teoh, P. Winkel, A. Anderson, J. Chen, J. Curtis, S. J. de Graaf, et al., Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements, Preprint at <http://arxiv.org/abs/2307.03169> (2023).
- [261] K. Sahay, J. Jin, J. Claes, J. D. Thompson, and S. Puri, High-Threshold Codes for Neutral-Atom Qubits with Biased Erasure Errors, *Physical Review X* **13**, 041013 (2023).
- [262] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature Communications* **9**, 4812 (2018).
- [263] K. Kumaran, M. Sajjan, S. Oh, and S. Kais, Random projection using random quantum circuits, *Physical Review Research* **6**, 013010 (2024).
- [264] S. Aaronson and S.-H. Hung, Certified Randomness from Quantum Supremacy, Preprint at <http://arxiv.org/abs/2303.01625> (2023).
- [265] S. Aaronson, Quantum Copy-Protection and Quantum Money, in *2009 24th Annual IEEE Conference on Computational Complexity* (2009), pp. 229–242, iSSN: 1093-0159.
- [266] N. Kumar, R. Mezher, and E. Kashefi, Efficient Construction of Quantum Physical Unclonable Functions with Unitary t-designs (2021).
- [267] F. G. S. L. Brandão, A. W. Harrow, and M. Horodecki, Local Random Quantum Circuits are Approximate Polynomial-Designs, *Communications in Mathematical Physics* **346**, 397 (2016).
- [268] G. Alagic, T. Gagliardoni, and C. Majenz, Unforgeable Quantum Encryption, in *Advances in Cryptology – EUROCRYPT 2018*, edited by J. B. Nielsen and V. Rijmen (Springer International Publishing, Cham, 2018), pp. 489–519, ISBN 978-3-319-78372-7.
- [269] A. Broadbent, C. E. González-Guillén, and C. Schuknecht, Quantum private broadcasting, *Physical Review A* **105**, 022606 (2022).

- [270] Z. Webb, The Clifford group forms a unitary 3-design, Preprint at <http://arxiv.org/abs/1510.02769> (2016).
- [271] H. Zhu, R. Kueng, M. Grassl, and D. Gross, The Clifford group fails gracefully to be a unitary 4-design, Preprint at <http://arxiv.org/abs/1609.08172> (2016).
- [272] B. Lévi, C. C. López, J. Emerson, and D. G. Cory, Efficient error characterization in quantum information processing, *Physical Review A* **75**, 022314 (2007).
- [273] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Symmetrized Characterization of Noisy Quantum Processes, *Science* **317**, 1893 (2007).
- [274] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, *Physical Review A* **80**, 12304 (2009).
- [275] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nature Physics* **14**, 595 (2018).
- [276] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, et al., A blueprint for demonstrating quantum supremacy with superconducting qubits, *Science* **360**, 195 (2018).
- [277] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Validating quantum computers using randomized model circuits, *Physical Review A* **100**, 32328 (2019).
- [278] F. G. Brandão, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill, Models of Quantum Complexity Growth, *PRX Quantum* **2**, 030316 (2021).
- [279] P. Hayden and J. Preskill, Black holes as mirrors: quantum information in random subsystems, *Journal of High Energy Physics* **2007**, 120 (2007).
- [280] L. Piroli, C. Sünderhauf, and X.-L. Qi, A random unitary circuit model for black hole evaporation, *Journal of High Energy Physics* **2020**, 63 (2020).
- [281] J. Emerson, Y. S. Weinstein, M. Saraceno, S. Lloyd, and D. G. Cory, Pseudo-Random Unitary Operators for Quantum Information Processing, *Science* **302**, 2098 (2003).
- [282] A. W. Harrow and R. A. Low, Random Quantum Circuits are Approximate 2-designs, *Communications in Mathematical Physics* **291**, 257 (2009).
- [283] M. Ohliger, V. Nesme, and J. Eisert, Efficient and feasible state tomography of quantum many-body systems, *New Journal of Physics* **15**, 15024 (2013).
- [284] Y. Nakata, C. Hirche, M. Koashi, and A. Winter, Efficient Quantum Pseudorandomness with Nearly Time-Independent Hamiltonian Dynamics, *Physical Review X* **7**, 21006 (2017).

- [285] A. W. Harrow, The Church of the Symmetric Subspace, Preprint at <http://arxiv.org/abs/1308.6595> (2013).
- [286] K.-M. Li, H. Dong, C. Song, and H. Wang, Approaching the chaotic regime with a fully connected superconducting quantum processor, *Physical Review A* **100**, 62302 (2019).
- [287] A. A. Mele, Introduction to Haar Measure Tools in Quantum Information: A Beginner's Tutorial, Preprint at <http://arxiv.org/abs/2307.08956> (2024).
- [288] G. E. P. Box and M. E. Muller, A Note on the Generation of Random Normal Deviates, *The Annals of Mathematical Statistics* **29**, 610 (1958).
- [289] C. E. Porter and R. G. Thomas, Fluctuations of Nuclear Reaction Widths, *Physical Review* **104**, 483 (1956).
- [290] G. Valiant and P. Valiant, Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs, in *Proceedings of the forty-third annual ACM symposium on Theory of computing* (Association for Computing Machinery, New York, NY, USA, 2011), STOC '11, pp. 685–694, ISBN 978-1-4503-0691-1.
- [291] L. M. Leemis and J. T. McQueston, Univariate Distribution Relationships, *The American Statistician* **62**, 45 (2008).
- [292] E. P. Wigner, in *The Collected Works of Eugene Paul Wigner: Part A: The Scientific Papers*, edited by A. S. Wightman (Springer, Berlin, Heidelberg, 1993), pp. 524–540, ISBN 978-3-662-02781-3.
- [293] F. J. Dyson, Statistical Theory of the Energy Levels of Complex Systems. I, *Journal of Mathematical Physics* **3**, 140 (1962).
- [294] T. Guhr, A. Müller–Groeling, and H. A. Weidenmüller, Random-matrix theories in quantum physics: common concepts, *Physics Reports* **299**, 189 (1998).
- [295] O. Bohigas, M. J. Giannoni, and C. Schmit, Characterization of Chaotic Quantum Spectra and Universality of Level Fluctuation Laws, *Physical Review Letters* **52**, 1 (1984).
- [296] J. W. Goodman, in *Laser Speckle and Related Phenomena*, edited by J. C. Dainty (Springer, Berlin, Heidelberg, 1975), pp. 9–75, ISBN 978-3-662-43205-1.
- [297] S. Mullane, Sampling random quantum circuits: a pedestrian's guide, Preprint at <http://arxiv.org/abs/2007.07872> (2020).
- [298] A. Ambainis and J. Emerson, Quantum t -designs: t -wise Independence in the Quantum World, in *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* (IEEE, 2007), pp. 129–140, ISBN 0-7695-2780-9, iSSN: 1093-0159.

- [299] I. Nechita and S. Singh, A graphical calculus for integration over random diagonal unitary matrices, *Linear Algebra and its Applications* **613**, 46 (2021).
- [300] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nature Physics* **16**, 1050 (2020).
- [301] P. Delsarte, J. M. Goethals, and J. J. Seidel, Spherical codes and designs, *Geometriae Dedicata* **6**, 363 (1977).
- [302] R. H. Hardin and N. J. A. Sloane, McLaren's Improved Snub Cube and Other New Spherical Designs in Three Dimensions, Preprint at <http://arxiv.org/abs/math/0207211> (2002).
- [303] M. Popp, F. Verstraete, M. A. Martín-Delgado, and J. I. Cirac, Localizable entanglement, *Physical Review A* **71**, 42306 (2005).
- [304] S. Goldstein, J. L. Lebowitz, C. Mastrodonato, R. Tumulka, and N. Zanghì, Universal Probability Distribution for the Wave Function of a Quantum System Entangled with its Environment, *Communications in Mathematical Physics* **342**, 965 (2016).
- [305] M. Rigol, V. Dunjko, and M. Olshanii, Thermalization and its mechanism for generic isolated quantum systems, *Nature* **452**, 854 (2008).
- [306] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, Quantum thermalization through entanglement in an isolated many-body system, *Science* **353**, 794 (2016).
- [307] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Reviews of Modern Physics* **91**, 21001 (2019).
- [308] M. Ueda, Quantum equilibration, thermalization and prethermalization in ultracold atoms, *Nature Reviews Physics* **2**, 669 (2020).
- [309] S. Popescu, A. J. Short, and A. Winter, Entanglement and the foundations of statistical mechanics, *Nature Physics* **2**, 754 (2006).
- [310] L. del Rio, A. Hutter, R. Renner, and S. Wehner, Relative thermalization, *Physical Review E* **94**, 22104 (2016).
- [311] R. Jozsa, D. Robb, and W. K. Wootters, Lower bound for accessible information in quantum mechanics, *Physical Review A* **49**, 668 (1994).
- [312] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì, On the Distribution of the Wave Function for Systems in Thermal Equilibrium, *Journal of Statistical Physics* **125**, 1193 (2006).
- [313] R. O'Donnell and J. Wright, Efficient Quantum Tomography, in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2016), pp. 899–912, ISBN 978-1-4503-4132-5, series Title: STOC '16.

- [314] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-Optimal Tomography of Quantum States, *IEEE Transactions on Information Theory* **63**, 5628 (2017).
- [315] W. K. Wootters, Statistical distance and Hilbert space, *Physical Review D* **23**, 357 (1981).
- [316] A. Morvan, B. Villalonga, X. Mi, S. Mandrà, A. Bengtsson, P. V. Klimov, Z. Chen, S. Hong, C. Erickson, I. K. Drozdov, et al., Phase transition in Random Circuit Sampling, Preprint at <http://arxiv.org/abs/2304.11119> (2023).
- [317] X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage, Preprint at <http://arxiv.org/abs/2112.01657> (2021).
- [318] B. Ware, A. Deshpande, D. Hangleiter, P. Niroula, B. Fefferman, A. V. Gorshkov, and M. J. Gullans, A sharp phase transition in linear cross-entropy benchmarking, Preprint at <http://arxiv.org/abs/2305.04954> (2023).
- [319] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Communications in Mathematical Physics* **28**, 251 (1972).
- [320] V. Khemani, A. Vishwanath, and D. A. Huse, Operator Spreading and the Emergence of Dissipative Hydrodynamics under Unitary Evolution with Conservation Laws, *Physical Review X* **8**, 31057 (2018).
- [321] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller, Cross-Platform Verification of Intermediate Scale Quantum Devices, *Physical Review Letters* **124**, 10504 (2020).
- [322] H. Kim, T. N. Ikeda, and D. A. Huse, Testing whether all eigenstates obey the eigenstate thermalization hypothesis, *Physical Review E* **90**, 52105 (2014).
- [323] K. Slagle, D. Aasen, H. Pichler, R. S. K. Mong, P. Fendley, X. Chen, M. Endres, and J. Alicea, Microscopic characterization of Ising conformal field theory in Rydberg chains, *Physical Review B* **104**, 235109 (2021).
- [324] D. N. Page, Average entropy of a subsystem, *Physical Review Letters* **71**, 1291 (1993).
- [325] M. Srednicki, Chaos and quantum thermalization, *Physical Review E* **50**, 888 (1994).
- [326] J. M. Deutsch, Quantum statistical mechanics in a closed system, *Physical Review A* **43**, 2046 (1991).
- [327] R. Nandkishore and D. A. Huse, Many-Body Localization and Thermalization in Quantum Statistical Mechanics, *Annual Review of Condensed Matter Physics* **6**, 15 (2015).
- [328] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, *Advances in Physics* **65**, 239 (2016).

- [329] J. M. Deutsch, Eigenstate thermalization hypothesis, *Reports on Progress in Physics* **81**, 082001 (2018).
- [330] G. Clos, D. Porras, U. Warring, and T. Schaetz, Time-Resolved Observation of Thermalization in an Isolated Quantum System, *Physical Review Letters* **117**, 170401 (2016).
- [331] F. Kranzl, A. Lasek, M. K. Joshi, A. Kalev, R. Blatt, C. F. Roos, and N. Yunger Halpern, Experimental Observation of Thermalization with Noncommuting Charges, *PRX Quantum* **4**, 020318 (2023).
- [332] H. Kim, Y. Park, K. Kim, H.-S. Sim, and J. Ahn, Detailed Balance of Thermalization Dynamics in Rydberg-Atom Quantum Simulators, *Physical Review Letters* **120**, 180502 (2018).
- [333] C. Neill, P. Roushan, M. Fang, Y. Chen, M. Kolodrubetz, Z. Chen, A. Megrant, R. Barends, B. Campbell, B. Chiaro, et al., Ergodic dynamics and thermalization in an isolated quantum system, *Nature Physics* **12**, 1037 (2016).
- [334] S. Pilatowsky-Cameo, C. B. Dag, W. W. Ho, and S. Choi, Complete Hilbert-Space Ergodicity in Quantum Dynamics of Generalized Fibonacci Drives, *Physical Review Letters* **131**, 250401 (2023).
- [335] A. J. Short, Equilibration of quantum systems and subsystems, *New Journal of Physics* **13**, 053009 (2011).
- [336] P. Reimann, Foundation of Statistical Mechanics under Experimentally Realistic Conditions, *Physical Review Letters* **101**, 190403 (2008).
- [337] M. Srednicki, The approach to thermal equilibrium in quantized chaotic systems, *Journal of Physics A: Mathematical and General* **32**, 1163 (1999).
- [338] C. Nation and D. Porras, Off-diagonal observable elements from random matrix theory: distributions, fluctuations, and eigenstate thermalization, *New Journal of Physics* **20**, 103003 (2018).
- [339] C. Nation and D. Porras, Ergodicity probes: using time-fluctuations to measure the Hilbert space dimension, *Quantum* **3**, 207 (2019).
- [340] M. Bauer, D. Bernard, and T. Jin, Universal fluctuations around typicality for quantum ergodic systems, *Physical Review E* **101**, 012115 (2020).
- [341] T. Kiendl and F. Marquardt, Many-Particle Dephasing after a Quench, *Physical Review Letters* **118**, 130601 (2017).
- [342] B. D. Clader, C. J. Trout, J. P. Barnes, K. Schultz, G. Quiroz, and P. Titum, Impact of correlations and heavy tails on quantum error correction, *Physical Review A* **103**, 052428 (2021).

- [343] C. Cafaro and S. Mancini, Quantum stabilizer codes for correlated and asymmetric depolarizing errors, *Physical Review A* **82**, 012306 (2010).
- [344] M. S. Carroll, J. R. Wootton, and A. W. Cross, Subsystem surface and compass code sensitivities to non-identical infidelity distributions on heavy-hex lattice, Preprint at <http://arxiv.org/abs/2402.08203> (2024).
- [345] Y. Nakata, P. S. Turner, and M. Muraō, Phase-random states: Ensembles of states with fixed amplitudes and uniformly distributed phases in a fixed basis, *Physical Review A* **86**, 012301 (2012).
- [346] P. Pfeuty, The one-dimensional Ising model with a transverse field, *Annals of Physics* **57**, 79 (1970).
- [347] S. Nadarajah, A Review of Results on Sums of Random Variables, *Acta Applicandae Mathematicae* **103**, 131 (2008).
- [348] A. M. Dalzell, N. Hunter-Jones, and F. G. S. L. Brandão, Random quantum circuits transform local noise into global white noise, Preprint at <http://arxiv.org/abs/2111.14907> (2021).
- [349] E. Paladino, Y. Galperin, G. Falci, and B. Altshuler, $1/f$ noise: Implications for solid-state quantum information, *Reviews of Modern Physics* **86**, 361 (2014).
- [350] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, Decoherence benchmarking of superconducting qubits, *npj Quantum Information* **5**, 1 (2019).
- [351] Y. Cai, Y. Tong, and J. Preskill, Stochastic error cancellation in analog quantum simulation, Preprint at <http://arxiv.org/abs/2311.14818> (2023).
- [352] A. J. Daley, Quantum trajectories and open many-body quantum systems, *Advances in Physics* **63**, 77 (2014).
- [353] J. Novak and M. LaCroix, Three lectures on free probability, Preprint at <http://arxiv.org/abs/1205.2097> (2012).
- [354] W. Dür, M. Hein, J. I. Cirac, and H.-J. Briegel, Standard forms of noisy quantum operations via depolarization, *Physical Review A* **72**, 052326 (2005).
- [355] N. Balakrishnan and V. B. Nevzorov, *A Primer on Statistical Distributions* (John Wiley & Sons, Ltd, 2003), ISBN 978-0-471-72222-9.
- [356] H. Shapourian, S. Liu, J. Kudler-Flam, and A. Vishwanath, Entanglement Negativity Spectrum of Random Mixed States: A Diagrammatic Approach, *PRX Quantum* **2**, 030347 (2021).
- [357] M. Žnidarič, Entanglement of random vectors, *Journal of Physics A: Mathematical and Theoretical* **40**, F105 (2006).

- [358] NIST digital library of mathematical functions.
- [359] C. L. Canonne, *A Survey on Distribution Testing: Your Data is Big. But is it Blue?*, no. 9 in Graduate Surveys (Theory of Computing Library, 2020).
- [360] R. W. Hoerl, Ridge Regression: A Historical Context, *Technometrics* **62**, 420 (2020).
- [361] F. Verstraete and H. Verschelde, On quantum channels, Preprint at <http://arxiv.org/abs/quant-ph/0202124> (2003).
- [362] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, Estimating the coherence of noise, *New Journal of Physics* **17**, 113020 (2015).
- [363] J. Acharya, A. Orlitsky, A. T. Suresh, and H. Tyagi, Estimating Renyi Entropy of Discrete Distributions, Preprint at <http://arxiv.org/abs/1408.1000> (2016).
- [364] J. Cotler, N. Hunter-Jones, and D. Ranard, Fluctuations of subsystem entropies at late times, arXiv:2010.11922v2 (2020).
- [365] M. Ippoliti and W. W. Ho, Dynamical Purification and the Emergence of Quantum State Designs from the Projected Ensemble, *PRX Quantum* **4**, 030322 (2023).
- [366] M. Ippoliti and W. W. Ho, Solvable model of deep thermalization with distinct design times, *Quantum* **6**, 886 (2022).
- [367] P. W. Claeys and A. Lamacraft, Emergent quantum state designs and biunitarity in dual-unitary circuit dynamics, *Quantum* **6**, 738 (2022).
- [368] M. Lucas, L. Piroli, J. De Nardis, and A. De Luca, Generalized deep thermalization for free fermions, *Physical Review A* **107**, 032215 (2023).
- [369] H. Wilming and I. Roth, High-temperature thermalization implies the emergence of quantum state designs, Preprint at <http://arxiv.org/abs/2202.01669> (2022).
- [370] T. Bhore, J.-Y. Desaulles, and Z. Papić, Deep thermalization in constrained quantum systems, *Physical Review B* **108**, 104317 (2023).
- [371] A. Chan and A. De Luca, Projected state ensemble of a generic model of many-body quantum chaos (2024).
- [372] N. D. Varikuti and S. Bandyopadhyay, Unraveling the emergence of quantum state designs in systems with symmetry (2024).
- [373] C. Vairogs and B. Yan, Extracting randomness from quantum 'magic' (2024).
- [374] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, et al., Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
- [375] X. Zhang, E. Kim, D. K. Mark, S. Choi, and O. Painter, A superconducting quantum simulator based on a photonic-bandgap metamaterial, *Science* **379**, 278 (2023).

- [376] Y. Liu, M. Otten, R. Bassirianjahromi, L. Jiang, and B. Fefferman, Benchmarking near-term quantum computers via random circuit sampling, Preprint at <http://arxiv.org/abs/2105.05232> (2022).
- [377] J.-S. Kim, L. S. Bishop, A. D. Córcoles, S. Merkel, J. A. Smolin, and S. Sheldon, Hardware-efficient random circuits to classify noise in a multiqubit system, *Physical Review A* **104**, 022609 (2021).
- [378] R. Harper, S. T. Flammia, and J. J. Wallman, Efficient learning of quantum noise, *Nature Physics* **16**, 1184 (2020).
- [379] C.-M. Chung, L. Bonnes, P. Chen, and A. M. Läuchli, Entanglement spectroscopy using quantum Monte Carlo, *Physical Review B* **89**, 195147 (2014).
- [380] H. F. Song, S. Rachel, C. Flindt, I. Klich, N. Laflorencie, and K. Le Hur, Bipartite fluctuations as a probe of many-body entanglement, *Physical Review B* **85**, 035409 (2012).
- [381] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, et al., Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, *Physical Review Letters* **127**, 180501 (2021).
- [382] Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, et al., Quantum computational advantage via 60-qubit 24-cycle random circuit sampling, *Science Bulletin* **67**, 240 (2022).
- [383] J. Preskill, Quantum computing and the entanglement frontier, Preprint at <http://arxiv.org/abs/1203.5813> (2012).
- [384] S. Ghosh, A. Deshpande, D. Hangleiter, A. V. Gorshkov, and B. Fefferman, Complexity Phase Transitions Generated by Entanglement, *Physical Review Letters* **131**, 30601 (2023).
- [385] K. Kechedzhi, S. V. Isakov, S. Mandrà, B. Villalonga, X. Mi, S. Boixo, and V. Smelyanskiy, Effective quantum volume, fidelity and computational cost of noisy quantum processing experiments, *Future Generation Computer Systems* **153**, 431 (2024).
- [386] J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, and J. Bermejo-Vega, Closing Gaps of a Quantum Advantage with Short-Time Hamiltonian Dynamics, *Physical Review Letters* **125**, 250501 (2020).
- [387] M. Mosca and A. Ekert, The Hidden Subgroup Problem and Eigenvalue Estimation on a Quantum Computer, Preprint at <http://arxiv.org/abs/quant-ph/9903071> (1999).
- [388] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, et al., Noisy intermediate-scale quantum algorithms, *Reviews of Modern Physics* **94**, 15004 (2022).

- [389] M. Hauru, A. Morningstar, J. Beall, M. Ganahl, A. Lewis, and G. Vidal, Simulation of quantum physics with Tensor Processing Units: brute-force computation of ground states and time evolution (2021).
- [390] T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, Measuring the capabilities of quantum computers, *Nature Physics* **18**, 75 (2022).
- [391] J. I. Cirac and P. Zoller, Goals and opportunities in quantum simulation, *Nature Physics* **8**, 264 (2012).
- [392] J. C. Bridgeman and C. T. Chubb, Hand-waving and interpretive dance: an introductory course on tensor networks, *Journal of Physics A: Mathematical and Theoretical* **50**, 223001 (2017).
- [393] G. Vidal, Efficient Classical Simulation of Slightly Entangled Quantum Computations, *Physical Review Letters* **91**, 147902 (2003).
- [394] Y. Zhou, E. M. Stoudenmire, and X. Waintal, What Limits the Simulation of Quantum Computers?, *Physical Review X* **10**, 41038 (2020).
- [395] S. Moses, C. Baldwin, M. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn, et al., A Race-Track Trapped-Ion Quantum Processor, *Physical Review X* **13**, 041052 (2023).
- [396] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, Quantum Algorithm for Simulating Real Time Evolution of Lattice Hamiltonians, *SIAM Journal on Computing* **0**, FOCS18 (2018).
- [397] M. A. Ganaie, M. Hu, A. K. Malik, M. Tanveer, and P. N. Suganthan, Ensemble deep learning: A review, *Engineering Applications of Artificial Intelligence* **115**, 105151 (2022).
- [398] A. J. Ferris and G. Vidal, Perfect sampling with unitary tensor networks, *Physical Review B* **85**, 165146 (2012).
- [399] T. Gorin, T. Prosen, T. H. Seligman, and M. Žnidarič, Dynamics of Loschmidt echoes and fidelity decay, *Physics Reports* **435**, 33 (2006).
- [400] W. H. Zurek, F. M. Cucchietti, and J. P. Paz, Gaussian Decoherence and Gaussian Echo from Spin Environments, Preprint at <http://arxiv.org/abs/quant-ph/0611200> (2007).
- [401] P. R. Zangara and H. M. Pastawski, Loschmidt echo in many-spin systems: a quest for intrinsic decoherence and emergent irreversibility, *Physica Scripta* **92**, 33001 (2017).
- [402] I. Najfeld and T. F. Havel, Derivatives of the matrix exponential and their computation, *Advances in applied mathematics* **16**, 321 (1995).
- [403] E. T. Bell, Exponential polynomials, *Annals of Mathematics* pp. 258–277 (1934).

- [404] C. S. Withers and S. Nadarajah, Moments from cumulants and vice versa, *International Journal of Mathematical Education in Science and Technology* **40**, 842 (2009).
- [405] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nature Physics* **14**, 447 (2018).
- [406] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nature Machine Intelligence* **1**, 155 (2019).
- [407] X. Zhang, M. Luo, Z. Wen, Q. Feng, S. Pang, W. Luo, and X. Zhou, Direct Fidelity Estimation of Quantum States Using Machine Learning, *Physical Review Letters* **127**, 130503 (2021).
- [408] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2023).
- [409] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, Measuring entanglement entropy in a quantum many-body system, *Nature* **528**, 77 (2015).
- [410] N. M. Linke, S. Johri, C. Figgatt, K. A. Landsman, A. Y. Matsuura, and C. Monroe, Measuring the Renyi entropy of a two-site Fermi-Hubbard model on a trapped ion quantum computer, *Physical Review A* **98**, 052334 (2018).
- [411] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, *Science* **364**, 260 (2019).
- [412] M. B. Plenio, Logarithmic Negativity: A Full Entanglement Monotone That is not Convex, *Physical Review Letters* **95**, 090503 (2005).
- [413] E. Bianchi, L. Hackl, M. Kieburg, M. Rigol, and L. Vidmar, Volume-Law Entanglement Entropy of Typical Pure Quantum States, *PRX Quantum* **3**, 030201 (2022).
- [414] A. Elben, R. Kueng, H.-Y. R. Huang, R. van Bijnen, C. Kokail, M. Dalmonte, P. Calabrese, B. Kraus, J. Preskill, P. Zoller, et al., Mixed-State Entanglement from Local Randomized Measurements, *Physical Review Letters* **125**, 200501 (2020).
- [415] G. J. Mooney, G. A. L. White, C. D. Hill, and L. C. L. Hollenberg, Whole-Device Entanglement in a 65-Qubit Superconducting Quantum Computer, *Advanced Quantum Technologies* **4**, 2100061 (2021).
- [416] S. Lee, D. P. Chi, S. D. Oh, and J. Kim, Convex-roof extended negativity as an entanglement measure for bipartite quantum systems, *Physical Review A* **68**, 062304 (2003).
- [417] U. T. Bhosale, S. Tomsovic, and A. Lakshminarayan, Entanglement between two subsystems, the Wigner semicircle and extreme-value statistics, *Physical Review A* **85**, 062331 (2012).

- [418] A. Datta, Negativity of random pure states, *Physical Review A* **81**, 052312 (2010).
- [419] M. Žnidarič, T. Prosen, G. Benenti, and G. Casati, Detecting entanglement of random states with an entanglement witness, *Journal of Physics A: Mathematical and Theoretical* **40**, 13787 (2007).
- [420] M. P. A. Fisher, V. Khemani, A. Nahum, and S. Vijay, Random Quantum Circuits, *Annual Review of Condensed Matter Physics* **14**, 335 (2023).
- [421] Y. Bao, S. Choi, and E. Altman, Theory of the phase transition in random unitary circuits with measurements, *Physical Review B* **101**, 104301 (2020).
- [422] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Reviews of Modern Physics* **82**, 277 (2010).
- [423] G. Vidal, Efficient Simulation of One-Dimensional Quantum Many-Body Systems, *Physical Review Letters* **93**, 040502 (2004).
- [424] K. Noh, L. Jiang, and B. Fefferman, Efficient classical simulation of noisy random quantum circuits in one dimension, *Quantum* **4**, 318 (2020).
- [425] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational Principle for Quantum Lattices, *Physical Review Letters* **107**, 70601 (2011).
- [426] I. L. Gutiérrez and C. B. Mendl, Real time evolution with neural-network quantum states, *Quantum* **6**, 627 (2022).
- [427] M. Ganahl, J. Beall, M. Hauru, A. G. M. Lewis, T. Wojno, J. H. Yoo, Y. Zou, and G. Vidal, Density Matrix Renormalization Group with Tensor Processing Units, *PRX Quantum* **4**, 10317 (2023).
- [428] T. Häner and D. S. Steiger, 0.5 Petabyte Simulation of a 45-Qubit Quantum Circuit, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Association for Computing Machinery, New York, NY, USA, 2017), ISBN 978-1-4503-5114-0, series Title: SC '17.
- [429] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, Locality and Digital Quantum Simulation of Power-Law Interactions, *Physical Review X* **9**, 031006 (2019).
- [430] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Physical Review B* **94**, 165116 (2016).
- [431] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).

- [432] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, et al., NetKet: A machine learning toolkit for many-body quantum systems, *SoftwareX* **10**, 100311 (2019).
- [433] D. Luo, Z. Chen, K. Hu, Z. Zhao, V. M. Hur, and B. K. Clark, Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models, *Physical Review Research* **5**, 013216 (2023).
- [434] A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman, An atomic boson sampler, Preprint at <http://arxiv.org/abs/2307.06936> (2023).
- [435] D. Zhu, Z. P. Cian, C. Noel, A. Risinger, D. Biswas, L. Egan, Y. Zhu, A. M. Green, C. H. Alderete, N. H. Nguyen, et al., Cross-platform comparison of arbitrary quantum states, *Nature Communications* **13**, 6620 (2022).
- [436] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, Theoretical and Experimental Perspectives of Quantum Verification, *PRX Quantum* **2**, 10102 (2021).
- [437] Y. Li and S. C. Benjamin, Efficient Variational Quantum Simulator Incorporating Active Error Minimization, *Physical Review X* **7**, 21050 (2017).
- [438] K. Temme, S. Bravyi, and J. M. Gambetta, Error Mitigation for Short-Depth Quantum Circuits, *Physical Review Letters* **119**, 180509 (2017).
- [439] Y. A. Lee and G. Vidal, Entanglement negativity and topological order, *Physical Review A* **88**, 42318 (2013).
- [440] T.-C. Lu, T. H. Hsieh, and T. Grover, Detecting Topological Order at Finite Temperature Using Entanglement Negativity, *Physical Review Letters* **125**, 116801 (2020).
- [441] S. Sang, Y. Li, T. Zhou, X. Chen, T. H. Hsieh, and M. P. A. Fisher, Entanglement Negativity at Measurement-Induced Criticality, *PRX Quantum* **2**, 30313 (2021).
- [442] S. Jandura, J. D. Thompson, and G. Pupillo, Optimizing Rydberg Gates for Logical-Qubit Performance, *PRX Quantum* **4**, 020336 (2023).
- [443] Y. Wang, K. Wang, E. F. Fenton, Y.-W. Lin, K.-K. Ni, and J. D. Hood, Reduction of laser intensity noise over 1 MHz band for single atom trapping, *Optics Express* **28**, 31209 (2020).
- [444] L. Li, W. Huie, N. Chen, B. DeMarco, and J. P. Covey, Active Cancellation of Servo-Induced Noise on Stabilized Lasers via Feedforward, *Physical Review Applied* **18**, 064005 (2022).
- [445] K.-N. Schymik, B. Ximenez, E. Bloch, D. Dreon, A. Signoles, F. Nogrette, D. Barredo, A. Browaeys, and T. Lahaye, In situ equalization of single-atom loading in large-scale optical tweezer arrays, *Physical Review A* **106**, 022611 (2022).

- [446] R. Tao, M. Ammenwerth, F. Gyger, I. Bloch, and J. Zeiher, High-fidelity detection of large-scale atom arrays in an optical lattice, Preprint at <http://arxiv.org/abs/2309.04717> (2023).
- [447] S. G. Menon, N. Glachman, M. Pompili, A. Dibos, and H. Bernien, An integrated atom array – nanophotonic chip platform with background-free imaging, Preprint at <http://arxiv.org/abs/2311.02153> (2023).
- [448] Y.-T. Chen, M. Szurek, B. Hu, J. d. Hond, B. Braverman, and V. Vuletić, High finesse bow-tie cavity for strong atom-photon coupling in Rydberg arrays, *Optics Express* **30**, 37426 (2022).
- [449] J. Ramette, J. Sinclair, Z. Vendeiro, A. Rudelis, M. Cetina, and V. Vuletić, Any-To-Any Connected Cavity-Mediated Architecture for Quantum Computing with Trapped Ions or Rydberg Arrays, *PRX Quantum* **3**, 010344 (2022).
- [450] J. P. Covey, H. Weinfurter, and H. Bernien, Quantum networks with neutral atom processing nodes, *npj Quantum Information* **9**, 1 (2023).
- [451] Y. Li and J. Thompson, High-rate and high-fidelity modular interconnects between neutral atom quantum processors, Preprint at <http://arxiv.org/abs/2401.04075> (2024).
- [452] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, et al., Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).
- [453] S. Julià-Farré, J. Vovrosh, and A. Dauphin, Amorphous quantum magnets in a two-dimensional Rydberg atom array, Preprint at <http://arxiv.org/abs/2402.02852> (2024).
- [454] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, et al., Suppressing quantum errors by scaling a surface code logical qubit, *Nature* **614**, 676 (2023).
- [455] Q. Xu, J. P. B. Ataiades, C. A. Pattison, N. Raveendran, D. Bluvstein, J. Wurtz, B. Vasic, M. D. Lukin, L. Jiang, and H. Zhou, Constant-Overhead Fault-Tolerant Quantum Computation with Reconfigurable Atom Arrays, Preprint at <http://arxiv.org/abs/2308.08648> (2023).
- [456] W. J. Eckner, A. W. Young, N. Schine, and A. M. Kaufman, High-power, fiber–laser-based source for magic-wavelength trapping in neutral-atom optical clocks, *Review of Scientific Instruments* **92**, 093001 (2021).
- [457] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman, Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice, *Science* **377**, 885 (2022).
- [458] T. Begušić, J. Gray, and G. K.-L. Chan, Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance, *Science Advances* **10**, eadk4321 (2024).

- [459] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, *Physical Review A* **74**, 022320 (2006).
- [460] F. Verstraete and J. I. Cirac, Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions, Preprint at <http://arxiv.org/abs/cond-mat/0407066> (2004).
- [461] G. Vidal, Class of Quantum Many-Body States That Can Be Efficiently Simulated, *Physical Review Letters* **101**, 110501 (2008).
- [462] V. Murg, F. Verstraete, and J. I. Cirac, Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states, *Physical Review A* **75**, 033605 (2007).
- [463] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, Architectures for Quantum Simulation Showing a Quantum Speedup, *Physical Review X* **8**, 021010 (2018).
- [464] S. Anand, C. E. Bradley, R. White, V. Ramesh, K. Singh, and H. Bernien, A dual-species Rydberg array, Preprint at <http://arxiv.org/abs/2401.10325> (2024).
- [465] F. Fang, K. Wang, V. S. Liu, Y. Wang, R. Cimmino, J. Wei, M. Bintz, A. Parr, J. Kemp, K.-K. Ni, et al., Probing critical phenomena in open quantum systems using atom arrays, Preprint at <http://arxiv.org/abs/2402.15376> (2024).
- [466] M. C. Tran, D. K. Mark, W. W. Ho, and S. Choi, Measuring Arbitrary Physical Properties in Analog Quantum Simulation, *Physical Review X* **13**, 011049 (2023).
- [467] K. Slagle, Y. Liu, D. Aasen, H. Pichler, R. S. K. Mong, X. Chen, M. Endres, and J. Alicea, Quantum spin liquids bootstrapped from Ising criticality in Rydberg arrays, *Physical Review B* **106**, 115122 (2022).
- [468] J. F. Wienand, S. Karch, A. Impertro, C. Schweizer, E. McCulloch, R. Vasseur, S. Gopalakrishnan, M. Aidelsburger, and I. Bloch, Emergence of fluctuating hydrodynamics in chaotic quantum systems, Preprint at <http://arxiv.org/abs/2306.11457> (2023).
- [469] P. T. Brown, D. Mitra, E. Guardado-Sanchez, R. Nourafkan, A. Reymbaut, C.-D. Hébert, S. Bergeron, A.-M. S. Tremblay, J. Kokalj, D. A. Huse, et al., Bad metallic transport in a cold atom Fermi-Hubbard system, *Science* **363**, 379 (2019).
- [470] V. B. Bulchandani and D. A. Huse, Hot band sound, Preprint at <http://arxiv.org/abs/2208.13767> (2022).
- [471] S. N. Thomas, B. Ware, J. D. Sau, and C. D. White, Comparing numerical methods for hydrodynamics in a 1D lattice model, Preprint at <http://arxiv.org/abs/2310.06886> (2023).

- [472] B. Bertini, F. Heidrich-Meisner, C. Karrasch, T. Prosen, R. Steinigeweg, and M. Žnidarič, Finite-temperature transport in one-dimensional quantum lattice models, *Reviews of Modern Physics* **93**, 025003 (2021).
- [473] T. Szołdra, P. Sierant, M. Lewenstein, and J. Zakrzewski, Unsupervised detection of decoupled subspaces: Many-body scars and beyond, *Physical Review B* **105**, 224205 (2022).
- [474] S. Moudgalya, B. A. Bernevig, and N. Regnault, Quantum many-body scars and Hilbert space fragmentation: a review of exact results, *Reports on Progress in Physics* **85**, 086501 (2022).
- [475] C. D. White, M. Zaletel, R. S. K. Mong, and G. Refael, Quantum dynamics of thermalizing systems, *Physical Review B* **97**, 035127 (2018).
- [476] X. Xu, S. Benjamin, J. Sun, X. Yuan, and P. Zhang, A Herculean task: Classical simulation of quantum computers, Preprint at <http://arxiv.org/abs/2302.08880> (2023).
- [477] A. Jamadagni, A. M. Läuchli, and C. Hempel, Benchmarking quantum computer simulation software packages, Preprint at <http://arxiv.org/abs/2401.09076> (2024).
- [478] A. W. Sandvik, Computational Studies of Quantum Spin Systems, *AIP Conference Proceedings* **1297**, 135 (2010).
- [479] R. Schmied, *Using Mathematica for Quantum Mechanics: A Student's Manual* (2020), arXiv:1403.7050.
- [480] A. Norambuena, D. Tancara, and R. Coto, Coding closed and open quantum systems in MATLAB: applications in quantum optics and condensed matter, *European Journal of Physics* **41**, 045404 (2020).
- [481] A. Weiße and H. Fehske, in *Computational Many-Particle Physics*, edited by H. Fehske, R. Schneider, and A. Weiße (Springer, Berlin, Heidelberg, 2008), pp. 529–544, ISBN 978-3-540-74686-7.
- [482] J.-H. Jung and J. D. Noh, Guide to Exact Diagonalization Study of Quantum Thermalization, *Journal of the Korean Physical Society* **76**, 670 (2020).
- [483] D. Jaksch, Numerical Methods for Strongly Correlated Systems, <https://exampleconference.files.wordpress.com/2015/10/nummethstrong.pdf> (2015).
- [484] F. Mezzadri, How to generate random matrices from the classical compact groups, Preprint at <http://arxiv.org/abs/math-ph/0609050> (2007).