

Statistical Mechanics of Problems in Transcription Regulation

Thesis by
Muir Morrison

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



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2021
Defended June 17, 2020

© 2021

Muir Morrison

ORCID: 0000-0002-0768-7234

All rights reserved

ACKNOWLEDGEMENTS

[Intentionally left blank]

ABSTRACT

As the quantity of sequenced genome data continues to multiply, our understanding of the transcriptional regulation of genomes has lagged behind. This deficit impinges on research throughout biology, from fundamental questions of how evolution proceeds to eminently practical questions such as how antibiotic resistance arises.

In this thesis we present three threads that address the question of transcriptional regulation from distinct perspectives. The first thread focuses on the simplest non-trivial regulation motif common in bacteria. We analyze in turn a sampling of the myriad mathematical models previously proposed in the literature for this system. We attempt to shine light on the similarities and differences of the models' predictions, clarify their microscopic interpretations, and offer guidance as to situations when one model or another should be preferred or even distinguishable.

The second thread considers a substantially more complicated genetic circuit, for which we build a minimal phenomenological model that retains intuitive microscopic meaning for all its parameters. The model neatly explains recent experimental observations of bistability in the circuit, and suggests natural generalizations to other metabolically important gene circuits with qualitatively similar architectures.

Motivation for the third thread comes from even more complicated transcriptional regulation problems with a multitude of regulatory proteins and binding sites, where even enumerating all possible DNA-protein complexes manually is a formidable challenge. Here we propose a method to tackle this complexity that uses ideas from quantum field theory to encode *assembly rules* for macromolecular complexes. By specifying a small set of rules, we avoid manual enumeration of the much larger set of complexes, allowing the formalism to automatically generate this set for us.

PUBLISHED CONTENT AND CONTRIBUTIONS

¹K. S. Laxhuber, M. J. Morrison, G. Chure, N. M. Belliveau, C. Strandkvist, K. L. Naughton, and R. Phillips, “Theoretical investigation of a genetic switch for metabolic adaptation”, PLOS ONE **15**, edited by L. M. Espinoza-Fonseca, e0226453 (2020). <https://doi.org/10.1371/journal.pone.0226453>

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	v
List of Illustrations	vii
List of Tables	xvii
Chapter I: Introduction	1
Chapter II: Reconciling Kinetic and Equilibrium Models of Bacterial Tran- scription	3
2.1 Abstract	3
2.2 Introduction	3
2.3 Mean Gene Expression	7
2.4 Beyond Means in Gene Expression	21
2.5 Finding the “right” model: Bayesian parameter inference	30
2.6 Discussion and future work	46
Chapter III: Theoretical Investigation of a Genetic Switch for Metabolic Adap- tation	56
3.1 Abstract	56
3.2 Introduction	56
3.3 Experimental motivation	59
3.4 Model	61
3.5 Results and discussion	68
3.6 Conclusion	81
Chapter IV: An efficient representation for statistical mechanics of multi- particle complexes in biology	88
4.1 Introduction	88
4.2 Formalism	91
Chapter A: Appendices for Chapter II, Reconciling Kinetic and Equilibrium Models of Bacterial Transcription	108
A.1 Derivations for non-bursty promoter models	108
A.2 Bursty promoter models - generating function solutions and numerics	129
A.3 Bayesian inference	147
Appendix B: Supporting information for Chapter III, Theoretical Investigation of a Genetic Switch for Metabolic Adaptation	163
B.1 Model choices, simplifications and assumptions	163
B.2 Parameter estimation	170
B.3 Additional plots and explanations of the results	175
B.4 Chemical master equation	178
B.5 Experimental materials and methods	181

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 An overview of the simple repression motif at the level of means.	
(A) Schematic of the qualitative biological picture of the simple repression genetic architecture. (B) and (C) A variety of possible mathematicized cartoons of simple repression, along with the effective parameter ρ which subsumes all regulatory details of the architecture that do not directly involve the repressor. (B) Simple repression models from an equilibrium perspective. (C) Equivalent models cast in chemical kinetics language. (D) The “master curve” to which all cartoons in (B) and (C) collapse.	6
2.2 Comparison of different models for noise in the constitutive promoter.	
(A) The left column depicts various plausible models for the dynamics of constitutive promoters. In model (1), transcripts are produced in a Poisson process [36, 57]. Model (2) features explicit modeling of RNAP binding/unbinding kinetics [58]. Model (3) is a more detailed generalization of model (2), treating transcription initiation as a multi-step process proceeding through closed and open complexes [38]. Model (4) is somewhat analogous to (2) except with the precise nature of active and inactive states left ambiguous [19, 23, 42]. Finally, model (5) can be viewed as a certain limit of model (4) in which transcripts are produced in bursts, and initiation of bursts is a Poisson process. The right column shows the Fano factor ν (variance/mean) for each model. Note especially the crucial diagnostic: (2) and (3) have ν strictly below 1, while only for (4) and (5) can ν exceed 1. Models with Fano factors ≤ 1 cannot produce the single-cell data observed in part (B) without introducing additional assumptions and model complexity. (B) Data from [36]. Mean mRNA count vs. Fano factor (variance/mean) for different promoters as determined with single-molecule mRNA Fluorescence <i>in situ</i> Hybridization. The colorbar indicates the predicted binding affinity of RNAP to the promoter sequence as determined in [54]. Numbers serve for cross comparison with data presented in Figure 3.	23

- 2.3 Constitutive promoter posterior inference and model comparison.** (A) The joint posterior density of model 5, the bursty promoter with negative binomially-distributed steady state, is plotted with MCMC samples. 1D marginal probability densities are plotted as flanking histograms. The model was fit on *lacUV5* data from [36]. (B) The empirical distribution function (ECDF) of the observed population distribution of mRNA transcripts under the control of a constitutive *lacUV5* promoter is shown in black. The median posterior predictive ECDFs for models (1), Poisson, and (5), negative binomial, are plotted in dark green and dark blue, respectively. Lighter green and blue regions enclose 95% of all posterior predictive samples from their respective models. Model (1) is in obvious contradiction with the data while model (5) is not. Single-cell mRNA count data is again from [36]. (C) Joint posterior distributions for burst rate k_i and mean burst size b for 18 unregulated promoters from [36]. Each contour indicates the 95% highest posterior probability density region for a particular promoter. Note that the vertical axis is shared with (D). (D) Plots of the burst rate k_i vs. the binding energy for each promoter as predicted in [54]. The dotted line shows the predicted slope according to Eq. 2.55, described in text. Each individual promoter is labeled with a unique number in both (C) and (D) for cross comparison and for comparison with Figure 2.2(B). 36

- 2.4 **Simple repression parameter inference and comparison.** (A) Contours which enclose 50% and 95% of the posterior probability mass are shown for each of several 2D slices of the 9D posterior distribution. The model assumes one unbinding rate for each operator (Oid, O1, O2) and one binding rate for each aTc induction concentration (corresponding to an unknown mean repressor copy number). (B, upper) Ratios of our inferred unbinding rates are compared with operator binding energy differences measured by Garcia and Phillips [33] (triangles) and Razo-Mejia et. al. [32] (circles). Blue glyphs compare O2-O1, while orange compare O1-Oid. Points with perfect agreement would lie on the dotted line. (B, lower) Unbinding rates for O1 (cyan) and Oid (red) inferred in this work are compared with single-molecule measurements from Hammar et. al. [56]. We plot the comparison assuming illustrative mRNA lifetimes of $\gamma^{-1} = 3$ min (triangles) or $\gamma^{-1} = 5$ min (squares). Dotted line is as in upper panel. (C) Theory-experiment comparison are shown for each of the datasets used in the inference of the model in (A). Observed single-molecule mRNA counts data from [36] are plotted as black lines. The median of the randomly generated samples for each condition is plotted as a dark colored line. Lighter colored bands enclose 95% of all samples for a given operator/repressor copy number pair. The unregulated promoter, *lacUV5*, is shown with each as a reference. 40
- 3.1 **A schematic of different genetic switches.** (A) and (B) show the two most well-known genetic switches: (A) two mutual repressors and (B) a self-activating gene. In (C), a very much simplified version of the circuit that we investigate in this paper can be seen, where the similarity to the switch in (B) is clear. A complete version of the model circuit can be found in Fig 3.3. 57

- 3.2 **Experimental data on the *xap* circuit.** (A) The expression of the *xapAB* promoter was measured for different extracellular concentrations of xanthosine (vertical axis). The left panel shows the wild type circuit while the middle and right panels show the effect of deleting the genes *xapABR* and *xapAB*, respectively. The wild type circuit behaves like a switch. Note that the fluorescence scale of the middle panel is not comparable with the other two, and also that the chosen xanthosine concentrations are different. (B) shows the fold-change in gene expression upon addition of xanthosine for the wt promoter (left panel) and for the promoter with only the XapR binding site adjacent to the polymerase binding site (right panel). Note that the fold-change used here differs from fold-change in, e.g., [23–25], in that no subtraction of autofluorescence was performed, which is adequate for the qualitative comparison of these two promoters. 60
- 3.3 **Model of the *xapAB* circuit.** The XapR dimers are induced by xanthosine and the induced XapR binds cooperatively as an activator to the *xapAB* promoter. For these two steps, quasi-equilibrium is assumed. If both XapR binding sites are occupied and the polymerase is bound, the gene is transcribed at rate r_m . The mRNA decays at rate γ_m , and both proteins are translated at rate r_p and decay at rate γ_p . XapA degrades xanthosine with Michaelis-Menten parameters k_a and K_a . Similarly, XapB is treated as a Michaelis-Menten enzyme which imports ($k_{b,i}$, $K_{b,i}$) and exports ($k_{b,e}$, $K_{b,e}$) xanthosine. Furthermore, xanthosine enters and leaves the cell through non-specific transport, proportional to rates k_{nup} and ξk_{nup} , respectively. 62
- 3.4 **The promoter states.** We consider only the completely occupied state as active and all other states (faded out in the figure) as completely inactive. The parameters are the interaction energy of the two XapR dimers ΔE_{coop} and the dissociation constants K_{XapR} and K_P of XapR and polymerase to the promoter, respectively. The concentrations of polymerase and active XapR are denoted by $[P]$ and $[XapR]$ 64

- 3.5 Phase portraits showing bistability.** 3D and 2D phase portraits for one set of parameters that leads to bistability. The parameter values are listed in Table 3.1. Note that all the concentrations ($[m]_a$, $[p]_a$, $[x]_a$) are measured in units of $K_a = 5 \cdot 10^4$ nM. The surfaces in (A) and the curves in (B) are the nullclines of the state variables, and their intersection points, marked in red in (B), are the steady-state solutions of the system. The region shaded in gray in (B) leads to negative concentrations and is unphysical. A vector plot of (B) that also shows the magnitude of flow at each point can be found in B.1. 69
- 3.6 Phase portraits for different extracellular xanthosine concentrations.** All parameters but $[c]_a$ are as presented in Table 3.1. The extracellular xanthosine concentration in these plots is $[c]_a = 7$ in (A) and $[c]_a = 40$ in (B) (recall that $[c]_a := \frac{c}{K_a}$ with $K_a = 5 \cdot 10^{-5}$ M, so $[c]_a$ is dimensionless). Tuning $[c]_a$ moves the orange line (xanthosine nullcline), but the blue curve (mRNA nullcline) is unchanged (see also B.1). It can clearly be seen that in (A) there is only the lower fixed point (fixed point number 1), whereas in (B) there is only the upper one (fixed point number 3). In between lies the bistable case that was shown in Fig 3.5. 72
- 3.7 Phase portraits without XapA or XapB.** All parameters are as presented in Table 3.1. In (A), the XapA term was removed from the kinetic equations. In (B), the equations lack the two terms from XapB. These plots clearly show that XapA has almost no influence on the qualitative behavior of the system (i.e. bistability and the order of magnitudes), but XapB is the essential feature for bistability. . . . 73

- 3.8 Phase portraits for less or no cooperativity.** Most parameters are as presented in Table 3.1, changes are mentioned below. Fixed points are marked in red. In (A), there is only one xanthosine binding site on XapR and everything unchanged for the XapR-promoter binding. Two parameters are changed: $\rho = 0.07$ and $[c]_a = 6$. This is necessary to compensate for the weaker induction such that the system is bistable. In (B), there is only one XapR binding site on the promoter and everything is unchanged for the xanthosine-XapR binding. Two parameters are changed: $\rho = 0.13$ and $[c]_a = 3$. In (C), there is only one xanthosine binding site on XapR and also only one XapR binding site on the the promoter. Two parameters are changed: $\rho = 0.1$ and $[XapR]_R = 5$. Whereas bistability is retained in (A) and (B), it cannot be obtained anymore in (C). 74
- 3.9 Distributions from stochastic simulations and the corresponding deterministic phase portraits.** Apart from $[c]_a$, the parameters are the same as in Table 3.1. The phase portraits were obtained from the deterministic system similarly to those presented in the previous sections. For the distributions, the simulations were run 5000 times for 10^6 s each (simulated time) and started at a mRNA, protein and intracellular xanthosine count of 0. We show the two cases of unimodality (low expression in (A) and high expression in (C)) as well as the case of bimodality in (B). The values of $[c]_a$ are 12 in (A), 18.5 in (B), and 25 in (C) (recall that $[c]_a := \frac{c}{K_a}$ with $K_a = 5 \cdot 10^{-5}$ M, so $[c]_a$ is dimensionless). The output from the stochastic simulations is in good agreement with the concentrations at the fixed points in the deterministic phase portraits. 77

- 3.10 Stochastic and deterministic time evolution of protein (XapA/XapB) and adaptation time.** (A) Shown in blue is the result from one typical run of the stochastic simulation, and in orange the trajectory obtained from solving the deterministic ODE's. In light blue, two more extreme runs of the simulation are shown for comparison. The simulation was run for $5 \cdot 10^5$ s each and started at an mRNA, protein and intracellular xanthosine count of 0. The parameters that were used are the same as in Table 3.1, the only exception being the extracellular xanthosine concentration, which was chosen to be $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M) just as in Figure 3.9(C). (B) In blue we plot the waiting time distribution for 1000 runs of the simulation (same conditions as (A)) to reach the steady-state mean. We define this time as the elapsed time when the protein copy number first reaches 90% of its value at the upper fixed point. To better visualize the bulk of the distribution, we excluded from the plot ~ 10 runs with adaptation times larger than $6 \cdot 10^5$ s. The green dashed line indicates the mean of the blue distribution. The orange line shows the corresponding deterministic time. 79
- 3.11 Bifurcation diagram showing the hysteresis loop in the stochastic system.** The blue line shows the positions of the deterministic fixed points, while the dashed line indicates the instability of the middle one. The orange and yellow points show the mean of the stochastic simulations when started at zero and at the high fixed point, respectively. The positions of the two peaks in the bimodal distributions are indicated by smaller points, connected by dotted lines. To make this clearer, the bimodal distributions themselves are shown on the right. The arrows illustrate the hysteresis loop in the stochastic simulations. 81
- A.1 Two-state promoter chemical master equation.** (A) Schematic of the two state promoter simple repression model. Rates k_R^+ and k_R^- are the association and dissociation rates of the transcriptional repressor, respectively, r is the transcription initiation rate, and γ is the mRNA degradation rate. (B) Schematic depiction of the mRNA count state transitions. The model in (A) only allows for jumps in mRNA of size 1. The production of mRNA can only occur when the promoter is in the transcriptionally active state. 108

- A.2 **Bursty transcription for unregulated promoter.** (A) Schematic of the one-state bursty transcription model. Rate k_i is the bursty initiation rate, γ is the mRNA degradation rate, and b is the mean burst size. (B) Schematic depiction of the mRNA count state transitions. The model in (A) allows for transitions of > 1 mRNA counts with probability $G_{m-m'}$, where the state jumps from having m' mRNA to having m mRNA in a single burst of gene expression. 129
- A.3 **Reindexing double sum.** Schematic for reindexing the sum $\sum_{m=0}^{\infty} \sum_{m'=0}^m$. Blue circles depict the 2D grid of nonnegative integers restricted to the lower triangular part of the m, m' plane. The trick is that this double sum runs over all (m, m') pairs with $m' \leq m$. Summing m first instead of m' requires determining the boundary: the upper boundary of the m' -first double sum becomes the lower boundary of the m -first double sum. 133
- A.4 **Schematic of three-stage promoter from [10].** Adapted from Shahrezaei & Swain [10]. In their paper they derive a closed form solution for the protein distribution. Our two-state bursty promoter at the mRNA level can be mapped into their solution with some relabeling. 138
- A.5 **Theory-data comparison of inference on unregulated promoters.** Comparison of the inference (red shaded area) vs the experimental measurements (black lines) for 18 different unregulated promoters with different mean mRNA expression levels from Ref. [8]. Upper panels show the empirical cumulative distribution function (ECDF), while the lower panels show the differences with respect to the median of the posterior samples. White numbers are the same as in Figure 2.1 for cross comparison. The predicted binding energies $\beta\Delta\epsilon_p$ were obtained from the energy matrix model in Ref. [20] 158
- A.6 **aTc controlled expression of LacI repressor.** Schematic of the circuit used in [8] to control the expression of the LacI repressor. The *lacI* gene is under the control of the TetR repressor. As the TetR repressor is inactivated upon binding of anhydrotetracycline or aTc, the more aTc added to the media were cells are growing, the less TetR repressors available to control the expression of the *lacI* gene, resulting in more LacI repressors per cell. LacI simultaneously controls the expression of the mRNA on which single-molecule mRNA FISH was performed for gene expression quantification. 159

- B.1 Results for different models of transcription.** The parameters are the same as the ones that were used for all other phase portraits in this report and that are listed in a table in the main part. In (A), the phase portrait after addition of non-zero transcription rates to all other polymerase-bound states is shown (third assumption lifted). The new parameter values are $r_I = 10^{-2}r_p$ for the transcription rate when only one XapR is bound and $r_0 = 10^{-3}r_p$ for the basal rate. In (B), the phase portrait after introducing different XapR dissociation constants can be seen (second assumption lifted). The new parameter values are $[XapR]_{R,1} = 10$ and $[XapR]_{R,2} = 0.1$ 164
- B.2 Kinetic scheme for transport by proton-driven pumps.** The mechanism of lac permease as for example described in [4] is used as a model for xanthosine transport. The enzyme binds a proton and a substrate (steps 1 and 2), undergoes conformational change (steps 3 and 4), and then releases substrate and proton on the other side of the membrane (steps 5 and 6). When the enzyme is empty, it can change its conformation again to repeat the transport (steps 7 and 8). Instead, it could also bind a new proton and substrate and transport them in the opposite direction (steps 6, 5, 4, 3, 2, 1). Alternatively, it could also keep the proton and just cycle back and forth, exchanging substrates from outside and inside the cell (steps 5, 4, 3, 2). 168
- B.3 Induction curves given by MWC model.** The main plot shows the probability of XapR being in the active state as a function of the logarithm of the xanthosine concentration for $\Delta\epsilon_x$ as shown, $K_{\chi A} = 10^2$, and $K_{IA} = 10^2$. The shaded areas indicate where the curve lies for changes in K_{IA} of one order of magnitude. Changing $K_{\chi A}$ by one order of magnitude, on the other hand, shifts all the curves by one order of magnitude in the corresponding direction to the left or right. The inset shows how the allowed regimes of $\Delta\epsilon_x$ and K_{IA} depend on each other. This explains why for $\Delta\epsilon_x = 7$ and $K_{IA} = 10$ (lower green line in the main plot), p_{active} at infinite xanthosine concentrations becomes far too low. 172
- B.4 Vector plot of a standard case of bistability.** The 2D phase portraits as a vector plot for for the standard set of parameters that was used in the main text. 176

- B.5 A family of curves for increasing $[c]_a$.** The blue, dashed curve is the protein nullcline, which does not change with $[c]_a$. The other curves are the xanthosine nullclines for different values of the extracellular xanthosine concentration. 176
- B.6 Time evolution and adaptation times from the stochastic simulation when transcriptional bursting is included.** (A) shows the time evolution from one run of the simulation. The simulation was run once for $5 \cdot 10^5$ s and started at an mRNA, protein and intracellular xanthosine count of 0. The parameters that were used are the same as in the table in the main text, the only exception being the extracellular xanthosine concentration, which was chosen to be $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M). (B) Blue shows the time until 90% of the protein concentration at the upper fixed point was reached in 1000 runs of the simulation (same conditions as (A)). The distribution from the main text (without bursting) is shown in orange for comparison. For both, the mean of the distribution is indicated by the line in the corresponding color. 178
- B.7 Distributions from the stochastic simulation when transcriptional bursting is included.** Blue shows the distributions from the simulation with bursts, orange shows those without bursts from the main text for comparison. The following values of $[c]_a$ were used: (A) $[c]_a = 12$, (B) $[c]_a = 17.5$ (not shown in main text), (C) $[c]_a = 18.5$, and (D) $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M). All other parameters are as in the table in the main text, and the simulation was run 1000 times for 10^6 s each (simulated time) and started at a mRNA, protein and intracellular xanthosine count of 0. 179
- B.8 Decreasing $[c]_a$ in the fully induced system.** Apart from $[c]_a$, the parameters are the same as in the table in the main text. For the distributions, the simulations were run 1000 times for 10^6 s each and started at the mRNA, protein and intracellular xanthosine counts of the high fixed point in the corresponding phase portraits. The values of $[c]_a$ are 12 in (A), 10 in (B), and 9 in (C). When comparing to the distributions in the main text where the simulation had initial concentrations of 0, a clear hysteresis effect can be observed. 180

LIST OF TABLES

<i>Number</i>	<i>Page</i>
3.1 Nondimensional parameters and their estimated values.	67
B.1 States and weights for a more precise model of transcription. . . .	165
B.2 Primers for deletion strains. Lower case sequences are homologous to a chromosomal locus. Upper case indicates sequences used for PCR amplification.	182

Chapter 1

INTRODUCTION

This thesis is a collection of three loosely connected projects, with the loose tethers being the methods and approaches of statistical mechanics applied to problems in, or problems motivated by, transcription regulation. Here we offer a brief outline of the chapters to follow.

Chapter 2 begins by considering what has been called “the hydrogen atom” of gene regulation, the simple repression motif. A multitude of quantitative models have been proposed in the literature for treating this system, but there has been little consensus as to which microscopic features are essential for these models to include and which features are irrelevant. We systematically and quantitatively dissect a selection of these models and establish which experimental observables can and cannot distinguish “right” from “wrong” models. We explain the basic reason for the robust prior success of the so-called thermodynamic models of transcription regulation. Having established the thermodynamic models’ utility, we verify that their model parameters are physically meaningful quantities and not mere fitting parameters using a self-consistency criterion we derive that all the models under consideration must satisfy.

Chapter 3 steps up the complexity of the transcriptional regulation problem. If we continue the analogy to atomic structure calculations, and simple repression in Chapter 2 is “the hydrogen atom,” the gene circuit considered in Chapter 3 is perhaps at the complexity of “aluminum.” It features a substantially more complex regulatory architecture with multiple competing feedback pathways. Despite the complexity and dearth of experimental data on the system, we were able to construct a surprisingly robust coarse-grained model of the essential features of the feedback circuit. The model neatly explained several features of this particular experimental system, and we believe it will be more widely applicable to many important metabolic pathways. This work was primarily carried out by a very talented visiting undergraduate student, Kathrin Laxhuber, who I had the pleasure of mentoring on this project. Most of the ideas arose collaboratively and organically during daily brainstorming over the course of a few months, and I supervised and provided feedback through many iterations of the writing process.

Chapter 4, though inspired by problems in transcription regulation, takes a more abstract approach. In the previous two chapters, merely enumerating the states available to the various molecular players was not a difficult task. As the number of available species, DNA binding sites, and pairwise interactions grow, even this simple enumeration quickly becomes challenging. How can we address this problem? The animating idea for this chapter is that this is “just” a combinatorics problem, and in fact is closely analagous to the combinatorics that arise in quantum field theory. We show in this chapter that the analogy can be made formal and useful: we can use the operator algebra from quantum field theory to encode *assembly rules* from which the theory naturally generates all possible complexes, rather than manually specifying each and every *complete* complex of macromolecules. For this snapshot of this work-in-progress, we show the basic formalism and demonstrate that it behaves as intuitively expected, allowing us to design assembly rules that generate the desired complexes for a model while excluding any nonphysical complexes.

*Chapter 2***RECONCILING KINETIC AND EQUILIBRIUM MODELS OF
BACTERIAL TRANSCRIPTION****2.1 Abstract**

The study of transcription remains one of the centerpieces of modern biology with implications in settings from development to metabolism to evolution to disease. Precision measurements using a host of different techniques including fluorescence and sequencing readouts have raised the bar for what it means to quantitatively understand transcriptional regulation. In particular our understanding of the simplest genetic circuit is sufficiently refined both experimentally and theoretically that it has become possible to carefully discriminate between different conceptual pictures of how this regulatory system works. This regulatory motif, originally posited by Jacob and Monod in the 1960s, consists of a single transcriptional repressor binding to a promoter site and inhibiting transcription. In this paper, we show how seven distinct models of this so-called simple-repression motif, based both on equilibrium and kinetic thinking, can be used to derive the predicted levels of gene expression and shed light on the often surprising past success of the equilibrium models. These different models are then invoked to confront a variety of different data on mean, variance and full gene expression distributions, illustrating the extent to which such models can and cannot be distinguished, and suggesting a two-state model with a distribution of burst sizes as the most potent of the seven for describing the simple-repression motif.

2.2 Introduction

Gene expression presides over much of the most important dynamism of living organisms. The level of expression of batteries of different genes is altered as a result of spatiotemporal cues that integrate chemical, mechanical and other types of signals. The original repressor-operator model conceived by Jacob and Monod in the context of bacterial metabolism has now been transformed into the much broader subject of gene regulatory networks in living organisms of all kinds [1–3]. One of the remaining outstanding challenges to have emerged in the genomic era is our continued inability to predict the regulatory consequences of different regulatory architectures, i.e. the arrangement and affinity of binding sites for transcription

factors and RNA polymerases on the DNA. This challenge stems first and foremost from our ignorance about what those architectures even are, with more than 60% of the genes even in an ostensibly well understood organism such as *E. coli* having no regulatory insights at all [4–7]. But even once we have established the identity of key transcription factors and their binding sites of a given promoter architecture, there remains the predictive challenge of understanding its input-output properties, an objective that can be met by a myriad of approaches using the tools of statistical physics [8–25]. One route to such predictive understanding is to focus on the simplest regulatory architecture and to push the theory-experiment dialogue as far and as hard as it can be pushed [26, 27]. If we demonstrate that we can pass that test by successfully predicting both the means and variance in gene expression at the mRNA level, then that provides a more solid foundation upon which to launch into more complex problems - for instance, some of the previously unknown architectures uncovered in [5] and [28].

To that end, in this paper we examine a wide variety of distinct models for the simple repression regulatory architecture. This genetic architecture consists of a DNA promoter regulated by a transcriptional repressor that binds to a single binding site as developed in pioneering early work on the quantitative dissection of transcription [29, 30]. All of the proposed models coarse-grain away some of the important microscopic features of this architecture that have been elucidated by generations of geneticists, molecular biologists and biochemists. One goal in exploring such coarse-grainings is to build towards the future models of regulatory response that will be able to serve the powerful predictive role needed to take synthetic biology from a brilliant exercise in enlightened empiricism to a rational design framework as in any other branch of engineering. More precisely, we want phenomenology in the sense of coarse-graining away atomistic detail, but still retaining biophysical meaning. For example, we are not satisfied with the strictly phenomenological approach offered by the commonly used Hill functions. As argued in [31], Hill functions are ubiquitous precisely because they coarse-grain away all biophysical details into inscrutable parameters. Studies like [32] have demonstrated that Hill functions are clearly insufficient since each new situation requires a completely new set of parameters. Such work requires a quantitative theory of how biophysical changes at the molecular level propagate to input-output functions at the genetic circuit level. In particular a key question is: at this level of coarse-graining, what microscopic details do we need to explicitly model, and how do we figure that out? For example, do we need to worry about all or even any of the steps that

individual RNA polymerases go through each time they make a transcript? Turning the question around, can we see any imprint of those processes in the available data? If the answer is no, then those processes are irrelevant for our purposes. Forward modeling and inverse (statistical inferential) modeling are necessary to tackle such questions.

Figure 2.1(A) shows the qualitative picture of simple repression that is implicit in the repressor-operator model. An operator, the binding site on the DNA for a repressor protein, may be found occupied by a repressor, in which case transcription is blocked from occurring. Alternatively, that binding site may be found unoccupied, in which case RNA polymerase (RNAP) may bind and transcription can proceed. The key assumption we make in this simplest incarnation of the repressor-operator model is that binding of repressor and RNAP in the promoter region of interest is exclusive, meaning that one or the other may bind, but never may both be simultaneously bound. It is often imagined that when the repressor is bound to its operator, RNAP is sterically blocked from binding to its promoter sequence. Current evidence suggests this is sometimes, but not always the case, and it remains an interesting open question precisely how a repressor bound far upstream is able to repress transcription [4]. Suggestions include “action-at-a-distance” mediated by kinks in the DNA, formed when the repressor is bound, that prevent RNAP binding. Nevertheless, our modeling in this work is sufficiently coarse-grained that we simply assume exclusive binding and leave explicit accounting of these details out of the problem.

The logic of the remainder of the paper is as follows. In section 2.3, we show how both thermodynamic models and kinetic models based upon the chemical master equation all culminate in the same underlying functional form for the fold-change in the average level of gene expression as shown in Figure 2.1(D). Section 2.4 goes beyond an analysis of the mean gene expression by asking how the same models presented in Figure 2.1(C) can be used to explore noise in gene expression. To make contact with experiment, all of these models must make a commitment to some numerical values for the key parameters found in each such model. Therefore in Section 2.5 we explore the use of Bayesian inference to establish these parameters and to rigorously answer the question of how to discriminate between the different models.

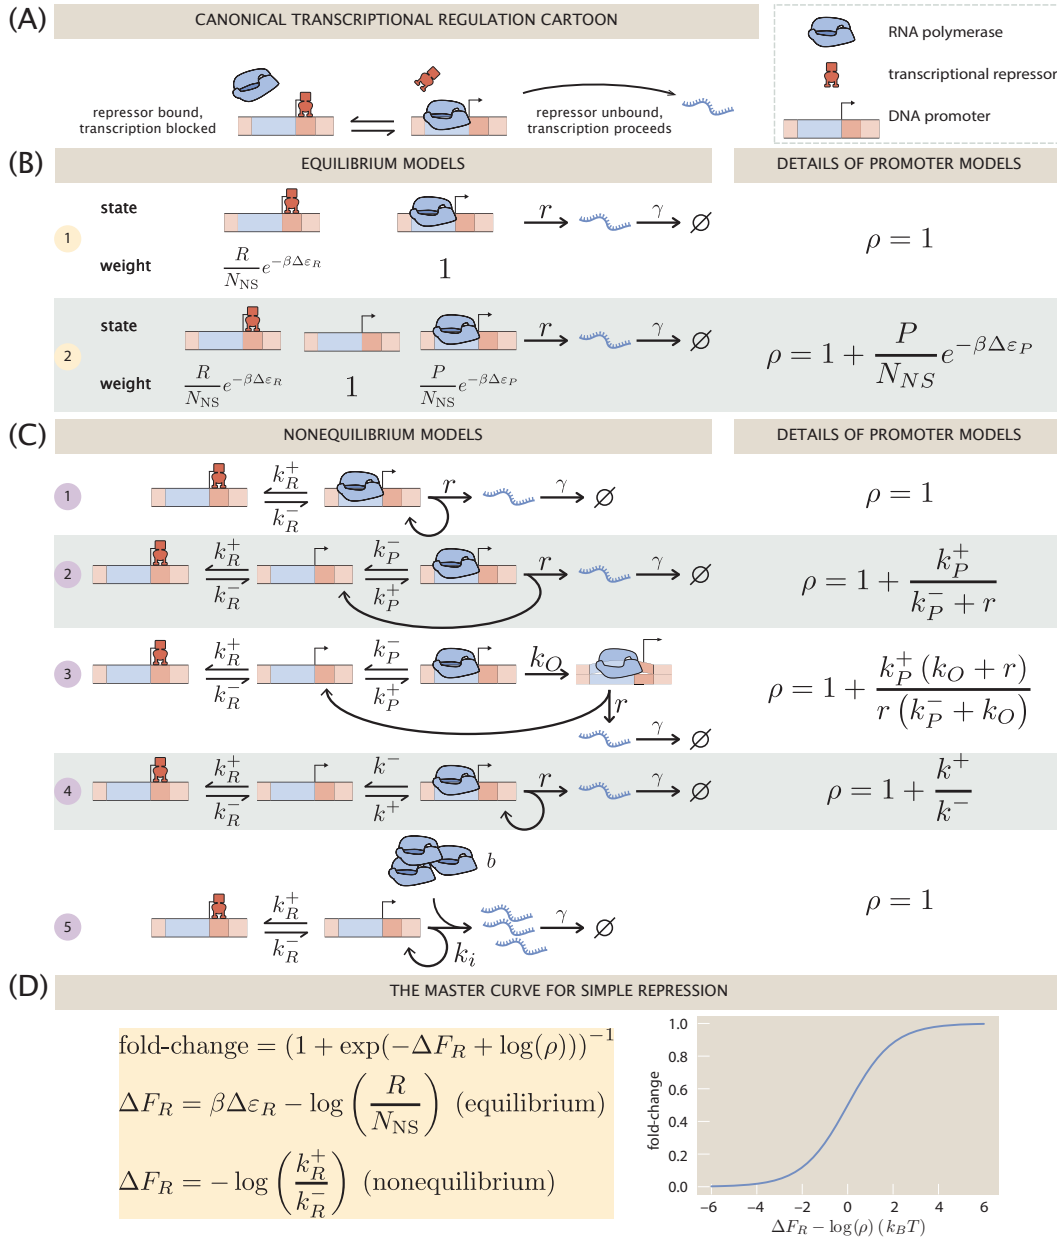


Figure 2.1: An overview of the simple repression motif at the level of means. (A) Schematic of the qualitative biological picture of the simple repression genetic architecture. (B) and (C) A variety of possible mathematicized cartoons of simple repression, along with the effective parameter ρ which subsumes all regulatory details of the architecture that do not directly involve the repressor. (B) Simple repression models from an equilibrium perspective. (C) Equivalent models cast in chemical kinetics language. (D) The “master curve” to which all cartoons in (B) and (C) collapse.

2.3 Mean Gene Expression

As noted in the previous section, there are two broad classes of models in play for computing the input-output functions of regulatory architectures as shown in Figure 2.1. In both classes of model, the promoter is imagined to exist in a discrete set of states of occupancy, with each such state of occupancy accorded its own rate of transcription – including no transcription for many of these states. The models are probabilistic with each state assigned some probability and the overall rate of transcription given by

$$\text{average rate of transcription} = \sum_i r_i p_i, \quad (2.1)$$

where i labels the distinct states, p_i is the probability of the i^{th} state, and r_i is the rate of transcription of that state. Ultimately, the different models differ along several key axes: what states to consider and how to compute the probabilities of those states.

The first class of models that are the focus of the present section on predicting the mean level of gene expression, sometimes known as thermodynamic models, invoke the tools of equilibrium statistical mechanics to compute the probabilities [8–17]. As seen in Figure 2.1(B), even within the class of thermodynamic models, we can make different commitments about the underlying microscopic states of the promoter. Indeed, the list of options considered here does not at all exhaust the suite of different microscopic states we can assign to the promoter.

The second class of models that allow us to access the mean gene expression use chemical master equations to compute the probabilities of the different microscopic states [18–25]. As seen in Figure 2.1(C), we consider a host of different nonequilibrium models, each of which will have its own result for both the mean (this section) and noise (next section) in gene expression.

Fold-changes are indistinguishable across models

As a first stop on our search for the “right” model of simple repression, let us consider what we can learn from theory and experimental measurements on the average level of gene expression in a population of cells. One experimental strategy that has been particularly useful (if incomplete since it misses out on gene expression dynamics) is to measure the fold-change in mean expression. The fold-change is defined as

$$\text{fold-change} = \frac{\langle \text{gene expression with repressor present} \rangle}{\langle \text{gene expression with repressor absent} \rangle} = \frac{\langle m(R > 0) \rangle}{\langle m(R = 0) \rangle} \quad (2.2)$$

where angle brackets $\langle \cdot \rangle$ denote the average over a population of cells and mean mRNA $\langle m \rangle$ is viewed as a function of repressor copy number R . What this means is that the fold-change in gene expression is a relative measurement of the effect of the transcriptional repressor ($R > 0$) on the gene expression level compared to an unregulated promoter ($R = 0$). The second equality in Eq. 2.2 follows from assuming that the translation efficiency, i.e., the number of proteins translated per mRNA, is the same in both conditions. In other words, we assume that mean protein level is proportional to mean mRNA level, and that the proportionality constant is the same in both conditions and therefore cancels out in the ratio. This is reasonable since the cells in the two conditions are identical except for the presence of the transcription factor, and the model assumes that the transcription factor has no direct effect on translation.

Fold-change has proven a very convenient observable in past work [32–35]. Part of its utility in dissecting transcriptional regulation is its ratiometric nature, which removes many secondary effects that are present when making an absolute gene expression measurement. Also, by measuring otherwise identical cells with and without a transcription factor present, any biological noise common to both conditions can be made to cancel away.

Figure 2.1 depicts a smorgasbord of mathematicized cartoons for simple repression using both thermodynamic and kinetic models that have appeared in previous literature. For each cartoon, we calculate the fold-change in mean gene expression as predicted by that model, deferring some algebraic details to Appendix A.1. What we will find is that all cartoons collapse to a single master curve, shown in Figure 2.1(D), which contains just two parameters. We label the parameters ΔF_R , an effective free energy parametrizing the repressor-DNA interaction, and ρ , which subsumes all details of transcription in the absence of repressors. We will offer some intuition for why this master curve exists and discuss why at the level of the mean expression, we are unable to discriminate “right” from “wrong” cartoons given only measurements of fold-changes in expression.

The Two-State Equilibrium Model

In this simplest model, depicted as (1) in Figure 2.1(B), the promoter is idealized as existing in one of two states, either repressor bound or repressor unbound. The rate of transcription is assumed to be proportional to the fraction of time spent in the repressor unbound state. From the relative statistical weights listed in Figure 2.1,

the probability p_U of being in the unbound state is

$$p_U = \left(1 + \frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R}\right)^{-1}. \quad (2.3)$$

The mean rate of transcription is then given by $r p_U$, as assumed by Eq. 2.1. The mean number of mRNA is set by the balance of average mRNA transcription and degradation rates, so it follows that the mean mRNA level is given by

$$\langle m \rangle = \frac{r}{\gamma} \left(1 + \frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R}\right)^{-1}, \quad (2.4)$$

where r is the transcription rate from the repressor unbound state, γ is the mRNA degradation rate, R is repressor copy number, N_{NS} is the number of nonspecific binding sites in the genome where repressors spend most of their time when not bound to the operator, $\beta \equiv 1/k_B T$, and $\Delta \varepsilon_R$ is the binding energy of a repressor to its operator site. The derivation of this result is deferred to Appendix A.1.

The fold-change is found as the ratio of mean mRNA with and without repressor as introduced in Eq. 2.2. Invoking that definition results in

$$\text{fold-change} = \left(1 + \frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R}\right)^{-1}, \quad (2.5)$$

which matches the form of the master curve in Figure 2.1(D) with $\rho = 1$ and $\Delta F_R = \beta \Delta \varepsilon_r - \log(R/N_{NS})$.

In fact it was noted in [35] that this two-state model can be viewed as the coarse-graining of any equilibrium promoter model in which no transcriptionally active states have transcription factor bound, or put differently, when there is no overlap between transcription factor bound states and transcriptionally active states. We will see this explicitly in the 3-state equilibrium model below, but perhaps surprising is that an analogous result carries over even to the nonequilibrium models we consider later.

The Three-State Equilibrium Model

Compared to the previous model, here we fine-grain the repressor unbound state into two separate states: empty, and RNAP bound as shown in (2) in Figure 2.1(B). This picture was used in [33] as we use it here, and in [32] and [35] it was generalized to incorporate small-molecule inducers that bind the repressor. The effect of this generalization is, roughly speaking, simply to rescale R from the total number of

repressors to a smaller effective number of available repressors which are unbound by inducers. We point out that the same generalization can be incorporated quite easily into any of our models in Figure 2.1 by simply rescaling the repressor copy number R in the equilibrium models, or equivalently k_R^+ in the nonequilibrium models.

The mean mRNA copy number, as derived in Appendix A.1 from a similar enumeration of states and weights as the previous model, is

$$\langle m \rangle = \frac{r}{\gamma} \frac{\frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}}{1 + \frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R} + \frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}}, \quad (2.6)$$

where the new variables are $\Delta \varepsilon_P$, the difference in RNAP binding energy to its specific site (the promoter) relative to an average nonspecific background site, and the RNAP copy number, P . The fold-change again follows immediately as

$$\text{fold-change} = \frac{\frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}}{1 + \frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R} + \frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}} \frac{1 + \frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}}{\frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}} \quad (2.7)$$

$$= \left(1 + \frac{\frac{R}{N_{NS}} e^{-\beta \Delta \varepsilon_R}}{1 + \frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}} \right)^{-1} \quad (2.8)$$

$$= (1 + \exp(-\Delta F_R - \log \rho))^{-1}, \quad (2.9)$$

with $\Delta F_R = \beta \Delta \varepsilon_R - \log(R/N_{NS})$ and $\rho = 1 + \frac{P}{N_{NS}} e^{-\beta \Delta \varepsilon_P}$ as shown in Figure 2.1(B). Thus far, we see that the two thermodynamic models, despite making different coarse-graining commitments, result in the same functional form for the fold-change in mean gene expression. We now explore how kinetic models fare when faced with computing the same observable.

The Poisson Promoter Nonequilibrium Model

For our first kinetic model, we imitate the states considered in the Two-State Equilibrium Model and consider the simplest possible picture with only two states, repressor bound and unbound. This is exactly the model used for the main results of [36]. In this picture, repressor association and dissociation rates from its operator site, k_R^+ and k_R^- , respectively, govern transitions between the two states. When the system is in the unbound state, transcription initiates at rate r , which represents a coarse-graining of all the downstream processes into a single effective rate. mRNA is degraded at rate γ as already exploited in the previous models.

Let $p_R(m, t)$ denote the joint probability of finding the system in the repressor bound state R with m mRNA molecules present at time t . Similarly define $p_U(m, t)$ for the repressor unbound state U . This model is governed by coupled master equations giving the time evolution of $p_R(m, t)$ and $p_U(m, t)$ [22, 24, 27] which we can write as

$$\begin{aligned}
\frac{d}{dt}p_R(m, t) &= - \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} + \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_R(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma p_R(m, t)}^{m \rightarrow m-1} \\
\frac{d}{dt}p_U(m, t) &= \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} - \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{r p_U(m-1, t)}^{m-1 \rightarrow m} - \overbrace{r p_U(m, t)}^{m \rightarrow m+1} \\
&\quad + \overbrace{(m+1)\gamma p_U(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma p_U(m, t)}^{m \rightarrow m-1},
\end{aligned} \tag{2.10}$$

where each term on the right corresponds to a transition between two states of the promoter as indicated by the overbrace label. In each equation, the first two terms describe transitions between promoter states due to repressors unbinding and binding, respectively. The final two terms describe degradation of mRNA, decreasing the copy number by one, and the terms with coefficient r describe transcription initiation increasing the mRNA copy number by one. We direct the reader to Appendix A.1 for a careful treatment showing how the form of this master equation follows from the corresponding cartoon in Figure 2.1.

We can greatly simplify the notation, which will be especially useful for the more complicated models yet to come, by re-expressing the master equation in vector form [37]. The promoter states are collected into a vector and the rate constants are collected into matrices as

$$\vec{p}(m) = \begin{pmatrix} p_R(m) \\ p_U(m) \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ \\ k_R^- & -k_R^+ \end{pmatrix}, \quad \mathbf{R} = \begin{pmatrix} 0 & 0 \\ 0 & r \end{pmatrix}, \tag{2.11}$$

so that the master equation may be condensed as

$$\frac{d}{dt}\vec{p}(m, t) = (\mathbf{K} - \mathbf{R} - \gamma m \mathbf{I}) \vec{p}(m, t) + \mathbf{R} \vec{p}(m-1, t) + \gamma(m+1) \mathbf{I} \vec{p}(m+1, t), \tag{2.12}$$

where \mathbf{I} is the identity matrix. Taking steady state by setting time derivatives to zero, the mean mRNA can be found to be

$$\langle m \rangle = \frac{r}{\gamma} \left(1 + \frac{k_R^+}{k_R^-} \right)^{-1}, \tag{2.13}$$

with the algebra details again deferred to Appendix A.1. Recall k_R^+ is proportional to the repressor copy number, so in computing fold-change, absence of repressor corresponds to $k_R^+ \rightarrow 0$. Therefore fold-change in this model is simply

$$\text{fold-change} = \left(1 + \frac{k_R^+}{k_R^-}\right)^{-1}, \quad (2.14)$$

again matching the master curve of Figure 2.1(D) with $\rho = 1$.

Nonequilibrium Model Two - RNAP Bound and Unbound States

Our second kinetic model depicted in Figure 2.1(C) mirrors the second equilibrium model of Figure 2.1(B) by fine-graining the repressor unbound state of nonequilibrium model 1, resolving it into an empty promoter state and an RNAP-bound state. Note in this picture, in contrast with model 4 below, transcription initiation is accompanied by a promoter state change, in keeping with the interpretation as RNAP-bound and empty states: if an RNAP successfully escapes the promoter and proceeds to elongation of a transcript, clearly it is no longer bound at the promoter. Therefore another RNAP must bind before another transcript can be initiated.

The master equation governing this model is analogous to Eqs. 2.11-2.12 for model 1 above. The main subtlety arises since transcription initiation accompanies a promoter state change. This can be understood by analogy to \mathbf{K} . The off-diagonal and diagonal elements of \mathbf{K} correspond to transitions arriving at or departing from, respectively, the promoter state of interest. If transcription initiation is accompanied by promoter state changes, we must have separate matrices for arriving and departing transcription events since the arriving and departing transitions have different initial copy numbers of mRNA, unlike for \mathbf{K} where they are the same (see Appendix A.1). The master equation for this model is

$$\frac{d}{dt}\vec{p}(m, t) = (\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m, t) + \mathbf{R}_A \vec{p}(m-1, t) + \gamma(m+1) \mathbf{I} \vec{p}(m+1, t), \quad (2.15)$$

with the state vector and promoter transition matrix defined as

$$\vec{p}(m) = \begin{pmatrix} p_R(m) \\ p_E(m) \\ p_P(m) \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- \\ 0 & k_P^+ & -k_P^- \end{pmatrix}, \quad (2.16)$$

and the initiation matrices given by

$$\mathbf{R}_A = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & r \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & r \end{pmatrix}. \quad (2.17)$$

The elements of $\vec{p}(m)$ encode the probabilities of having m mRNA present along with the promoter having repressor bound (R), being empty (E), or having RNAP bound (P), respectively. \mathbf{R}_A describes probability flux *arriving* at the state $\vec{p}(m)$ from a state with one fewer mRNA, namely $\vec{p}(m-1)$, and \mathbf{R}_D describes probability flux *departing* from the state $\vec{p}(m)$ for a state with one more mRNA, namely $\vec{p}(m+1)$. \mathbf{K} is closely analogous to model 1.

Mean mRNA at steady state is found analogously to model 1, with the result

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_R^- k_P^+}{k_R^- k_P^+ + k_R^- (k_P^- + r) + k_R^+ (k_P^- + r)}, \quad (2.18)$$

and with details again deferred to Appendix A.1. Fold-change is again found from the ratio prescribed by Eq. 2.2, from which we have

$$\text{fold-change} = \frac{k_R^- k_P^+}{k_R^- k_P^+ + k_R^- (k_P^- + r) + k_R^+ (k_P^- + r)} \frac{k_P^+ + k_P^- + r}{k_P^+} \quad (2.19)$$

$$= \left(1 + \frac{k_R^+}{k_R^-} \frac{k_P^- + r}{k_P^+ + k_P^- + r} \right)^{-1} \quad (2.20)$$

$$= \left(1 + \frac{k_R^+}{k_R^-} \left(1 + \frac{k_P^+}{k_P^- + r} \right)^{-1} \right)^{-1}, \quad (2.21)$$

which follows the master curve of Figure 2.1(D) with $\rho = 1 + k_P^+ / (k_P^- + r)$ as claimed.

Nonequilibrium Model Three - Multistep Transcription Initiation and Escape

One might reasonably complain that the first two “nonequilibrium” models we have considered are straw men. Their steady states necessarily satisfy detailed balance which is equivalent to thermodynamic equilibrium. Why is this the case? At steady state there is by definition no net probability flux in or out of each promoter state, but since the promoter states form a linear chain, there is only one way in or out of the repressor bound and RNAP bound states, implying each edge must actually have a net zero probability flux, which is the definition of detailed balance (usually phrased as equality of forward and reverse transition fluxes).

Now we consider model 3 in Figure 2.1(C) which allows the possibility of true nonequilibrium steady-state fluxes through the promoter states. We point out that this model was considered previously in [38] where a comparison was made with model 1 as used in [36]. The authors of [38] argued that the additional complexity is essential to properly account for the noise in the mRNA distribution. We will weigh in on both models later when we consider observables beyond fold-change.

The master equation governing this model is identical in form to model 2 above, namely

$$\frac{d}{dt}\vec{p}(m, t) = (\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m, t) + \mathbf{R}_A \vec{p}(m-1, t) + \gamma(m+1) \mathbf{I} \vec{p}(m+1, t), \quad (2.22)$$

but with a higher-dimensional state space and different matrices. The state vector and promoter transition matrix are now

$$\vec{p}(m) = \begin{pmatrix} p_R(m) \\ p_E(m) \\ p_C(m) \\ p_O(m) \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ & 0 & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- & 0 \\ 0 & k_P^+ & -k_P^- - k_O & 0 \\ 0 & 0 & k_O & 0 \end{pmatrix}, \quad (2.23)$$

with the four promoter states, in order, being repressor bound (R), empty (E), RNAP closed complex (C), and RNAP open complex (O). Besides increasing dimension by one, the only new feature in \mathbf{K} is the rate k_O , representing the rate of open complex formation from the closed complex, which we assume for simplicity to be irreversible in keeping with some [38] but not all [39] past literature. The initiation matrices are given by

$$\mathbf{R}_A = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r \end{pmatrix}, \quad (2.24)$$

again closely analogous to nonequilibrium model 2.

The expression for mean mRNA is substantially more complicated now, as worked out in Appendix A.1 where we find

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_R^- k_P^+ k_O}{k_R^- [(k_P^+ (k_O + r) + r(k_P^- + k_O))] + k_R^+ r(k_P^- + k_O)}, \quad (2.25)$$

which can be simplified to

$$\langle m \rangle = \frac{r}{\gamma} \frac{\frac{k_P^+ k_O}{r(k_O + k_P^-)}}{1 + \frac{k_P^+ (k_O + r)}{r(k_O + k_P^-)} + \frac{k_R^+}{k_R^-}}. \quad (2.26)$$

The strategy is to isolate the terms involving the repressor, so that now the fold-change is seen to be simply

$$\text{fold-change} = \frac{\frac{k_P^+ k_O}{r(k_O + k_P^-)}}{1 + \frac{k_P^+(k_O + r)}{r(k_O + k_P^-)} + \frac{k_R^+}{k_R^-}} \frac{1 + \frac{k_P^+(k_O + r)}{r(k_O + k_P^-)}}{\frac{k_P^+ k_O}{r(k_O + k_P^-)}} \quad (2.27)$$

$$= \left(1 + \frac{k_R^+}{k_R^-} \left(1 + \frac{k_P^+(k_O + r)}{r(k_O + k_P^-)} \right)^{-1} \right)^{-1}, \quad (2.28)$$

surprisingly reducing to the master curve of Figure 2.1(D) once again, with $\rho = 1 + \frac{k_P^+(k_O + r)}{r(k_O + k_P^-)}$.

This example hints that an arbitrarily fine-grained model of downstream transcription steps may still be collapsed to the form of the master curve for the means given in Figure 2.1(D), so long as the repressor binding is exclusive with transcriptionally active states. We offer this as a conjecture, and we suspect that a careful argument using the King-Altman diagram method [40, 41] might furnish a “proof.” Our focus here is not on full generality but rather to survey an assortment of plausible models for simple repression that have been proposed in the literature.

Nonequilibrium Model Four - “Active” and “Inactive” States

Model 4 in Figure 2.1(C) is at the core of the theory in [42]. At a glance the cartoon for this model may appear very similar to model 2, and mathematically it is, but the interpretation is rather different. In model 2, we interpreted the third state literally as an RNAP-bound promoter and modeled initiation of a transcript as triggering a promoter state change, making the assumption that an RNAP can only make one transcript at a time. In contrast, in the present model the promoter state does *not* change when a transcript is initiated. So we no longer interpret these states as literally RNAP bound and unbound but instead as coarse-grained “active” and “inactive” states, the details of which we leave unspecified for now. We will comment more on this model below when we discuss Fano factors of models.

Mathematically this model is very similar to models 1 and 2. Like model 1, the matrix R describing transcription initiation is diagonal, namely

$$\mathbf{R} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & r \end{pmatrix}. \quad (2.29)$$

The master equation takes verbatim the same form as it did for model 1, Eq. 2.12. Meanwhile the promoter transition matrix \mathbf{K} is the same as Eq. 2.16 from model 2, although we relabel the rate constants from k_p^\pm to k^\pm to reiterate that these are not simply RNAP binding and unbinding rates.

Carrying out the algebra, the mean mRNA can be found to be

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_R^- k^+}{k_R^- k^+ + k_R^- k^- + k_R^+ k^-}, \quad (2.30)$$

and the fold-change readily follows,

$$\text{fold-change} = \frac{k_R^- k^+}{k_R^- k^+ + k_R^- k^- + k_R^+ k^-} \frac{k_R^- k^+ + k_R^- k^-}{k_R^- k^+} \quad (2.31)$$

$$= \left(1 + \frac{k_R^+}{k_R^-} \left(1 + \frac{k^+}{k^-} \right)^{-1} \right)^{-1}, \quad (2.32)$$

from which we see $\rho = 1 + k^+/k^-$ as shown in Figure 2.1(C).

Nonequilibrium Model Five - Bursty Promoter

The final model we consider shown in Figure 2.1(C) is an intuitive analog to model 1, with just two states, repressor bound or unbound, and transition rates between them of k_R^+ and k_R^- . In model 1, when in the unbound state, single mRNA transcripts are produced as a Poisson process with some characteristic rate r . The current model by contrast produces, at some Poisson rate k_i , *bursts* of mRNA transcripts. The burst sizes are assumed to be geometrically distributed with a mean burst size b , which we will motivate in Section 2.4 when we derive this model as a certain limiting case of model 4.

From this intuitive picture and by analogy to model 1, then, it should be plausible that the mean mRNA level is

$$\langle m \rangle = \frac{k_i b}{\gamma} \left(1 + \frac{k_R^+}{k_R^-} \right)^{-1}, \quad (2.33)$$

which will turn out to be correct from a careful calculation. For now, we simply note that just like model 1, the fold-change becomes

$$\text{fold-change} = \left(1 + \frac{k_R^+}{k_R^-} \right)^{-1} \quad (2.34)$$

with $\rho = 1$ also like model 1. We will also see later how this model emerges as a natural limit of model 4.

Discussion of Results Across Models for Fold-Changes in Mean Expression

The key outcome of our analysis of the models in Figure 2.1 is the existence of a master curve shown in Figure 2.1(D) to which the fold-change predictions of all the models collapse. This master curve is parametrized by only two effective parameters: ΔF_R , which characterizes the number of repressors and their binding strength to the DNA, and ρ , which characterizes all other features of the promoter architecture. The key assumption underpinning this result is that no transcription occurs when a repressor is bound to its operator. Note, however, that we are agnostic about the molecular mechanism which achieves this; steric effects are one plausible mechanism, but, for instance, “action-at-a-distance” mediated by kinked DNA due to repressors bound tens or hundreds of nucleotides upstream of a promoter is plausible as well.

Why does the master curve of Figure 2.1(D) exist at all? This brings to mind the deep questions posed in, e.g., [31] and [43], suggesting we consider multiple plausible models of a system and search for their common patterns to tease out which broad features are and are not important. In our case, the key feature seems to be the exclusive nature of repressor and RNAP binding, which allows the parameter describing the repressor, ΔF_R , to cleanly separate from all other details of the promoter architecture, which are encapsulated in ρ . Arbitrary nonequilibrium behavior can occur on the rest of the promoter state space, but it may all be swept up in the effective parameter ρ , to which the repressor makes no contribution. We point the interested reader to [44] and [45] for an interesting analysis of similar problems using a graph-theoretic language.

As suggested in [35], we believe this master curve should generalize to architectures with multiple repressor binding sites, as long as the exclusivity of transcription factor binding and transcription initiation is maintained. The interpretation of ΔF_R is then of an effective free energy of all repressor bound states. In an equilibrium picture this is simply given by the log of the sum of Boltzmann weights of all repressor bound states, which looks like the log of a partition function of a subsystem. In a nonequilibrium picture, while we can still mathematically gather terms and give the resulting collection the label ΔF_R , it is unclear if the physical interpretation as an effective free energy makes sense. The problem is that free energies cannot be assigned unambiguously to states out of equilibrium because the free energy change along a generic path traversing the state space is path dependent, unlike at equilibrium. A consequence of this is that, out of equilibrium, ΔF_R is no longer a

simple sum of Boltzmann weights. Instead it resembles a restricted sum of King-Altman diagrams [40, 41]. Following the work of Hill [46], it may yet be possible to interpret this expression as an effective free energy, but this remains unclear to us. We leave this an open problem for future work.

If we relax the requirement of exclusive repressor-RNAP binding, one could imagine models in which repressor and RNAP doubly-bound states are allowed, where the repressor’s effect is to reduce the transcription rate rather than setting it to zero. Our results do not strictly apply to such a model, although we note that if the repressor’s reduction of the transcription rate is substantial, such a model might still be well-approximated by one of the models in Figure 2.1.

One may worry that our “one curve to rule them all” is a mathematical tautology. In fact we *agree* with this criticism if ΔF_R is “just a fitting parameter” and cannot be meaningfully interpreted as a real, physical free energy. An analogy to Hill functions is appropriate here. One of their great strengths and weaknesses, depending on the use they are put to, is that their parameters coarse-grain many details and are generally not interpretable in terms of microscopic models, for deep reasons discussed at length in [31]. By contrast, our master curve claims to have the best of both worlds: a coarse-graining of all details besides the repressor into a single effective parameter ρ , while simultaneously retaining an interpretation of ΔF_R as a physically meaningful and interpretable free energy. Our task, then, is to prove or disprove this claim.

How do we test this and probe the theory with fold-change measurements? There is a fundamental limitation in that the master curve is essentially a one-parameter function of $\Delta F_R + \log \rho$. Worse, there are many *a priori* plausible microscopic mechanisms that could contribute to the value of ρ , such as RNAP binding and escape kinetics [38, 39], and/or supercoiling accumulation and release [47, 48], and/or, RNAP clusters analogous to those observed in eukaryotes [49, 50] and recently also observed in bacteria [51]. Even if ΔF_R is measured to high precision, inferring the potential microscopic contributions to ρ , buried inside a log no less, from fold-change measurements seems beyond reach. As a statistical inference problem it is entirely nonidentifiable, in the language of [52], Section 4.3.

If we cannot simply infer values of ρ from measurements of fold-change, can we perturb some of the parameters that make up ρ and measure the change? Unfortunately we suspect this is off-limits experimentally: most of the potential contributors to ρ are global processes that affect many or all genes. For instance, changing RNAP as-

sociation rates by changing RNAP copy numbers, or changing supercoiling kinetics by changing topoisomerase copy numbers, would massively perturb the entire cell's physiology and confound any determination of ρ .

One might instead imagine a bottom-up modeling approach, where we mathematize a model of what we hypothesize the important steps are and are not, use *in vitro* data for the steps deemed important, and *predict* what ρ should be. But again, because of the one-parameter nature of the master curve, many different models will likely make indistinguishable predictions, and without any way to experimentally perturb *in vivo*, there is no clear way to test whether the modeling assumptions are correct.

In light of this, we prefer the view that parameters and rates are not directly comparable between cartoons in Figure 2.1. Rather, parameters in the simpler cartoons represent coarse-grained combinations of parameters in the finer-grained models. For instance, by equating ρ between any two models, one can derive various possible correspondences between the two models' parameters. Note that these correspondences are clearly not unique, since many possible associations could be made. It then is a choice as to what microscopic interpretations the model-builder prefers for the parameters in a particular cartoon, and as to which coarse-grainings lend intuition and which seem nonsensical. Indeed, since it remains an open question what microscopic features dominate ρ (as suggested above, perhaps RNAP binding and escape kinetics [38, 39], or supercoiling accumulation and release [47, 48], or something more exotic like RNAP clusters [49–51]), we are hesitant to put too much weight on any one microscopic interpretation of model parameters that make up ρ .

One possible tuning knob to probe ρ that would not globally perturb the cell's physiology is to manipulate RNAP binding sites. Work such as [53] has shown that models of sequence-dependent RNAP affinity can be inferred from data, and the authors of [54] showed that the model of [53] has predictive power by using the model to *design* binding sites of a desired affinity. But for our purposes, this begs the question: the authors of [53] *assumed* a particular model (essentially our 3-state equilibrium model but without the repressor), so it is unclear how or if such methods can be turned around to *compare* different models of promoter function.

Another possible route to dissect transcription details without a global perturbation would be to use phage polymerase with phage-specific promoters. While such results would carry some caveats, e.g., whether the repression of the phage polymerase is a good analog to the repression of the native RNAP, it could nevertheless be worthy

of consideration.

We have already pointed out that the master curve of Figure 2.1 is essentially a one-parameter model, the one parameter being $\Delta F_R + \log \rho$. By now the reader may be alarmed as to how can we even determine ΔF_R and ρ independently of each other, never mind shedding a lens on the internal structure of ρ itself. A hint is provided by the weak promoter approximation, invoked repeatedly in prior studies [14, 32, 33] of simple repression using the 3-state equilibrium model in Figure 2.1(B). In that picture, the weak promoter approximation means $\frac{P}{N_{NS}} \exp(-\beta \Delta \epsilon_P) \ll 1$, meaning therefore $\rho \approx 1$. This approximation can be well justified on the basis of the number of RNAP and σ factors per cell and the strength of binding of RNAP to DNA at weak promoters. This is suggestive, but how can we be sure that ρ is not, for instance, actually 10^2 and that ΔF_R hides a compensatory factor? A resolution is offered by an independent inference of ρ in the absence of repressors. This was done in [42] by fitting nonequilibrium model 4 in Figure 2.1(C), with zero repressor (looking ahead, this is equivalent to model 4 in Figure 2.2(A)), to single-cell mRNA counts data from [34]. This provided a determination of k^+ and k^- , from which their ratio is estimated to be no more than a few 10^{-1} and possibly as small as 10^{-2} .

The realization that $\rho \approx 1$ to an excellent approximation, *independent* of which model in Figure 2.1 one prefers, goes a long way towards explaining the surprising success of equilibrium models of simple repression. Even though our 2- and 3-state models get so many details of transcription wrong, it does not matter because fold-change is a cleverly designed ratio. Since ρ subsumes all details except the repressor, and $\log \rho \approx 0$, fitting these simple models to fold-change measurements can still give a surprisingly good estimate of repressor binding energies. So the ratiometric construction of fold-change fulfills its intended purpose of canceling out all features of the promoter architecture except the repressor itself. Nevertheless it is perhaps surprising how effectively it does so: *a priori*, one might not have expected ρ to be quite so close to 1.

We would also like to highlight the relevance of [55] here. Landman et. al. reanalyzed and compared *in vivo* and *in vitro* data on the lacI repressor's binding affinity to its various operator sequences. (The *in vivo* data was from, essentially, fitting our master curve to expression measurements.) They find broad agreement between the *in vitro* and *in vivo* values. This reinforces the suspicion that the equilibrium $\Delta \epsilon_R$ repressor binding energies do in fact represent real physical free energies. Again, *a priori* this did not have to be the case, even knowing that $\rho \approx 1$.

In principle, if ΔF_R can be measured to sufficient precision, then deviations from $\rho = 1$ become a testable matter of experiment. In practice, it is probably unrealistic to measure repressor rates k_R^+ or k_R^- or fold-changes in expression levels (and hence $\Delta \varepsilon_R$) precisely enough to detect the expected tiny deviations from $\rho = 1$. We can estimate the requisite precision in ΔF_R to resolve a given $\Delta \rho$ by noting, since $\rho \approx 1$, that $\log(1 + \Delta \rho) \approx \Delta \rho$, so $\Delta(\Delta F_R) \approx \Delta \rho$. Suppose we are lucky and $\Delta \rho$ is ~ 0.1 , on the high end of our range estimated above. A determination of $\Delta \varepsilon_R/k_B T$ with an uncertainty of barely 0.1 was achieved in the excellent measurements of [32], so this requires a very difficult determination of ΔF_R for a very crude determination of ρ , which suggests, to put it lightly, this is not a promising path to pursue experimentally. It is doubtful that inference of repressor kinetic rates would be any easier.

Moving forward, we have weak evidence supporting the interpretation of ΔF_R as a physically real free energy [55] and other work casting doubt [56]. How might we resolve the confusion? If there is no discriminatory power to test the theory and distinguish the various models with measurements of fold-changes in means, how do we probe the theory? Clearly to discriminate between the nonequilibrium models in Figure 2.1, we need to go beyond means to ask questions about kinetics, noise and even full distributions of mRNA copy numbers over a population of cells. If the “one-curve-to-rule-them-all” is more than a mathematical tautology, then the free energy of repressor binding inferred from fold-change measurements should agree with repressor binding and unbinding rates. In other words, the equilibrium and nonequilibrium definitions of ΔF_R must agree, meaning

$$\Delta F_R = \beta \Delta \varepsilon_R - \log(R/N_{NS}) = -\log(k_R^+/k_R^-), \quad (2.35)$$

must hold, where $\beta \Delta \varepsilon_R$ is inferred from the master curve fit to fold-change measurements, and k_R^+ and k_R^- are inferred in some orthogonal manner. Single molecule measurements such as from [56] have directly observed these rates, and in the remainder of this work we explore a complementary approach: inferring repressor binding and unbinding rates k_R^+ and k_R^- from single-molecule measurements of mRNA population distributions.

2.4 Beyond Means in Gene Expression

In this section, our objective is to explore the same models considered in the previous section, but now with reference to the the question of how well they describe the distribution of gene expression levels, with special reference to the variance in these distributions. To that end, we repeat the same pattern as in the previous section by

examining the models one by one. In particular we will focus on the Fano factor, defined as the variance/mean. This metric serves as a powerful discriminatory tool from the null model that the steady-state mRNA distribution must be Poisson, giving a Fano factor of one.

Kinetic models for unregulated promoter noise

Before we can tackle simple repression, we need an adequate phenomenological model of constitutive expression. The literature abounds with options from which we can choose, and we show several potential kinetic models for constitutive promoters in Figure 2.2(A). Let us consider the suitability of each model for our purposes in turn.

Noise in the Poisson Promoter Model

The simplest model of constitutive expression that we can imagine is shown as model 1 in Figure 2.2(A) and assumes that transcripts are produced as a Poisson process from a single promoter state. This is the picture from Jones et. al. [36] that was used to interpret a systematic study of gene expression noise over a series of promoters designed to have different strengths. This model insists that the “true” steady-state mRNA distribution is Poisson, implying the Fano factor ν must be 1. In [36], the authors carefully attribute measured deviations from Fano = 1 to intensity variability in fluorescence measurements, gene copy number variation, and copy number fluctuations of the transcription machinery, e.g., RNAP itself. In this picture, the master equation makes no appearance, and all the corrections to Poisson behavior are derived as additive corrections to the Fano factor. For disproving the “universal noise curve” from So et. al. [59], this picture was excellent. It is appealing in its simplicity, with only two parameters, the initiation rate r and degradation rate γ . Since γ is independently known from other experiments, and the mean mRNA copy number is r/γ , r is easily inferred from data. In other words, the model is not excessively complex for the data at hand. But for many interesting questions, for instance in the recent work [42], knowledge of means and variances alone is insufficient, and a description of the full distribution of molecular counts is necessary. For this we need a (slightly) more complex model than model 1 that would allow us to incorporate the non-Poissonian features of constitutive promoters directly into a master equation formulation.

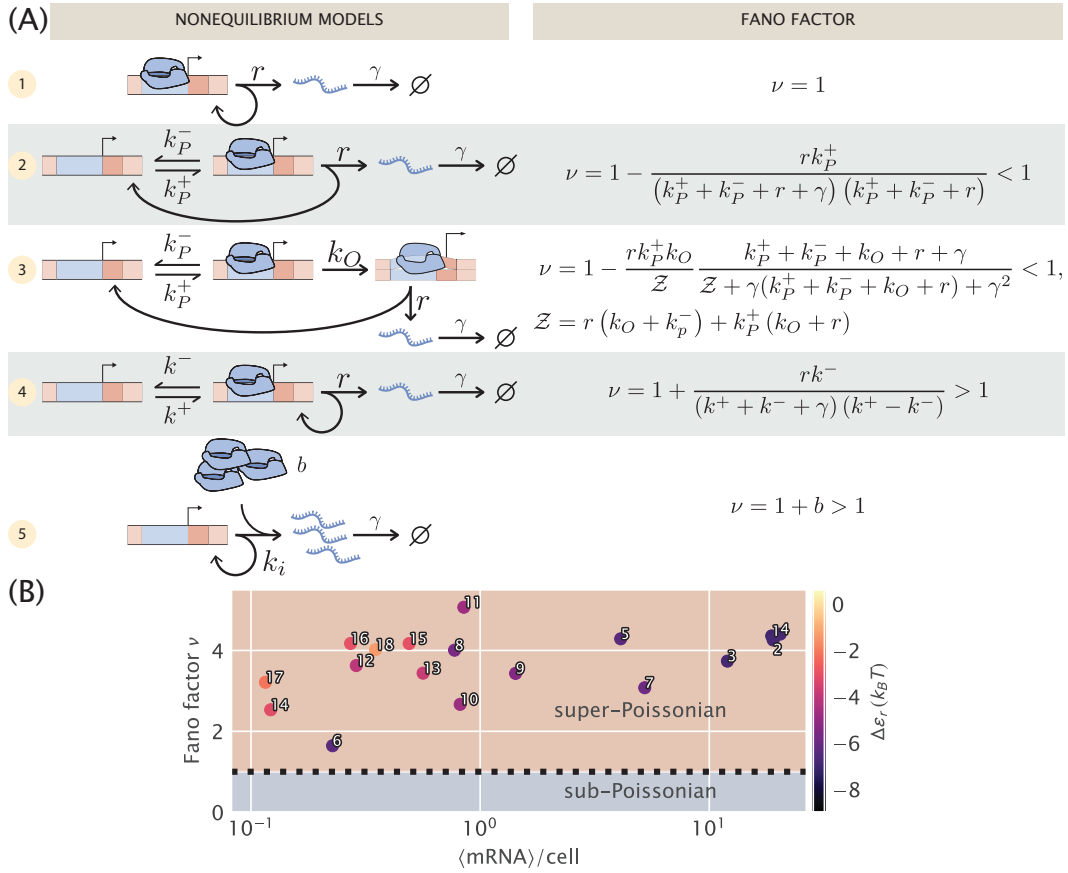


Figure 2.2: **Comparison of different models for noise in the constitutive promoter.** (A) The left column depicts various plausible models for the dynamics of constitutive promoters. In model (1), transcripts are produced in a Poisson process [36, 57]. Model (2) features explicit modeling of RNAP binding/unbinding kinetics [58]. Model (3) is a more detailed generalization of model (2), treating transcription initiation as a multi-step process proceeding through closed and open complexes [38]. Model (4) is somewhat analogous to (2) except with the precise nature of active and inactive states left ambiguous [19, 23, 42]. Finally, model (5) can be viewed as a certain limit of model (4) in which transcripts are produced in bursts, and initiation of bursts is a Poisson process. The right column shows the Fano factor ν (variance/mean) for each model. Note especially the crucial diagnostic: (2) and (3) have ν strictly below 1, while only for (4) and (5) can ν exceed 1. Models with Fano factors ≤ 1 cannot produce the single-cell data observed in part (B) without introducing additional assumptions and model complexity. (B) Data from [36]. Mean mRNA count vs. Fano factor (variance/mean) for different promoters as determined with single-molecule mRNA Fluorescence *in situ* Hybridization. The colorbar indicates the predicted binding affinity of RNAP to the promoter sequence as determined in [54]. Numbers serve for cross comparison with data presented in Figure 3.

Noise in the Two-State Promoter, RNAP Bound or Unbound.

Our second model of constitutive transcription posits an architecture in which the promoter is either empty or bound by RNAP [27, 58]. Here, as shown in model 2 of Figure 2.2(A), transcription initiation results in a state transition from the bound to the unbound state, reflecting the microscopic reality that an RNAP that has begun to elongate a transcript is no longer available at the start site to begin another. As shown in Appendix A.1, the Fano factor in this model is given by

$$\nu = 1 - \frac{rk_P^+}{(k_P^+ + k_P^- + r)(\gamma + k_P^+ + k_P^- + r)}. \quad (2.36)$$

The problem with this picture is that the Fano factor is always < 1 . To make contact with the experimental reality of $\nu > 1$ as shown in Figure 2.2(B), clearly some corrections will be needed. While this model adds an appealing element of microscopic reality, we are forced to reject it as the additional complexity is unable to capture the phenomenology of interest. Obviously the promoter state does in fact proceed through cycles of RNAP binding, initiating, and elongating, but it seems that the super-Poissonian noise in mRNA copy number we want to model must be governed by other features of the system.

Noise in the Three-State Promoter, Multistep Transcription Initiation and Escape.

How might we remedy the deficits of model 2? It is known [39] that once RNAP initially binds the promoter region, a multitude of distinct steps occur sequentially before RNAP finally escapes into the elongation phase. Perhaps adding some of this mechanistic detail as shown in model 3 of Figure 2.2(A) might rescue the previous model. The next simplest refinement of that model could consider open complex formation and promoter escape as separate steps rather than as a single effective step. In other words, we construct model 3 by adding a single extra state to model 2, and we will label the two RNAP-bound states as the closed and open complexes, despite the true biochemical details certainly being more complex. For example, earlier work extended this model by adding an additional repressor bound state and did not explicitly consider the limit with no repressor that we analyze here [38]. Again, our goal here is not a complete accounting of all the relevant biochemical detail; this is an exploratory search for the important features that a model needs to include to square with the known experimental reality of constitutive expression.

Unfortunately, as hinted at in earlier work [38], this model too has Fano factor $\nu < 1$. We again leave the algebraic details for Appendix A.1 and merely state the result that

$$\nu = 1 - \frac{rk_P^+k_O}{\mathcal{Z}} \frac{k_P^+ + k_P^- + k_O + r + \gamma}{\mathcal{Z} + \gamma(k_P^+ + k_P^- + k_O + r) + \gamma^2}, \quad (2.37)$$

where we defined $\mathcal{Z} = r(k_O + k_P^-) + k_P^+(k_O + r)$ for notational tidiness. This is necessarily less than 1 for arbitrary rate constants.

In fact, we suspect *any* model in which transcription proceeds through a multistep cycle must necessarily have $\nu < 1$. The intuitive argument compares the waiting time distribution to traverse the cycle with the waiting time for a Poisson promoter (model 1) with the same mean time. The latter is simply an exponential distribution. The former is a convolution of multiple exponentials, and intuitively the waiting time distribution for a multistep process should be more peaked with a smaller fractional width than a single exponential with the same mean. A less disperse waiting time distribution means transcription initiations are more uniformly distributed in time relative to a Poisson process. Hence the distribution of mRNA over a population of cells should be less variable compared to Poisson, giving $\nu < 1$. (In Appendix A.1 we present a more precise version of the intuitive arguments in this paragraph.) Regardless of the merits of this model in describing the noise properties of constitutive transcription initiation, it ultimately fails the simplest quantitative feature of the data, namely that the Fano factor > 1 and hence we must discard this mechanistic picture and search elsewhere.

Noise in a Two-State Promoter with “Active” and “Inactive” States

Inspired by [42], we next revisit an analog of model 2 in Figure 2.2(A), but as with the analogous models considered in Section 2.3, the interpretation of the two states is changed. Rather than explicitly viewing them as RNAP bound and unbound, we view them as “active” and “inactive,” which are able and unable to initiate transcripts, respectively. We are noncommittal as to the microscopic details of these states.

One interpretation [47, 48] for the active and inactive states is that they represent the promoter’s supercoiling state: transitions to the inactive state are caused by accumulation of positive supercoiling, which inhibits transcription, and transitions back to “active” are caused by gyrase or other topoisomerases relieving the supercoiling. This is an interesting possibility because it would mean the timescale for promoter

state transitions is driven by topoisomerase kinetics, not by RNAP kinetics. From in vitro measurements, the former are suggested to be of order minutes [47]. Contrast this with model 2, where the state transitions are assumed to be governed by RNAP, which, assuming a copy number per cell of order 10^3 , has a diffusion-limited association rate $k_{on} \sim 10^2 \text{ s}^{-1}$ to a target promoter. Combined with known K_d 's of order μM , this gives an RNAP dissociation rate k_{off} of order 10^2 s^{-1} . As we will show below, however, there are some lingering puzzles with interpreting this supercoiling hypothesis, so we leave it as a speculation and refrain from assigning definite physical meaning to the two states in this model.

Intuitively one might expect that, since transcripts are produced as a Poisson process only when the promoter is in one of the two states in this model, transcription initiations should now be “bunched” in time, in contrast to the “anti-bunching” of models 2 and 3 above. One might further guess that this bunching would lead to super-Poissonian noise in the mRNA distribution over a population of cells. Indeed, as shown in Appendix A.1, a calculation of the Fano factor produces

$$\nu = 1 + \frac{rk^-}{(k^+ + k^- + \gamma)(k^+ + k^-)}, \quad (2.38)$$

which is strictly greater than 1, verifying the above intuition. Note we have dropped the P label on the promoter switching rates to emphasize that these very likely do not represent kinetics of RNAP itself. This calculation can also be sidestepped by noting that the model is mathematically equivalent to the simple repression model from [36], with states and rates relabeled and reinterpreted.

How does this model compare to model 1 above? In model 1, all non-Poisson features of the mRNA distribution were handled as extrinsic corrections. By contrast, here the 3 parameter model is used to fit the full mRNA distribution as measured in mRNA FISH experiments. In essence, all variability in the mRNA distribution is regarded as “intrinsic,” arising either from stochastic initiation or from switching between the two coarse-grained promoter states. The advantage of this approach is that it fits neatly into the master equation picture, and the parameters thus inferred can be used as input for more complicated models with regulation by transcription factors.

While this seems promising, there is a major drawback for our purposes which was already uncovered by the authors of [42]: the statistical inference problem is nonidentifiable, in the sense described in Section 4.3 of [52]. What this means is that it is impossible to infer the parameters r and k^- from the single-cell mRNA

counts data of [36] (as shown in Fig. S2 of [42]). Rather, only the ratio r/k^- could be inferred. In that work, the problem was worked around with an informative prior on the ratio k^-/k^+ . That approach is unlikely to work here, as, recall, our entire goal in modeling constitutive expression is to use it as the basis for a yet more complicated model, when we add on repression. But adding more complexity to a model that is already poorly identified is a fool's errand, so we will explore one more potential model.

Noise Model for One-State Promoter with Explicit Bursts

The final model we consider is inspired by the failure mode of model 4. The key observation above was that, as found in [42], only two parameters, k^+ and the ratio r/k^- , could be directly inferred from the single-cell mRNA data from [36]. So let us take this seriously and imagine a model where these are the only two model parameters. What would this model look like?

To develop some intuition, consider model 4 in the limit $k^+ \ll k^- \lesssim r$, which is roughly satisfied by the parameters inferred in [42]. In this limit, the system spends the majority of its time in the inactive state, occasionally becoming active and making a burst of transcripts. This should call to mind the well-known phenomenon of transcriptional bursting, as reported in, e.g., [47, 48, 60]. Let us make this correspondence more precise. The mean dwell time in the active state is $1/k^-$. While in this state, transcripts are produced at a rate r per unit time. So on average, r/k^- transcripts are produced before the system switches to the inactive state. Once in the inactive state, the system dwells there for an average time $1/k^+$ before returning to the active state and repeating the process. r/k^- resembles an average burst size, and $1/k^+$ resembles the time interval between burst events. More precisely, $1/k^+$ is the mean time between the end of one burst and the start of the next, whereas $1/k^+ + 1/k^-$ would be the mean interval between the start of two successive burst events, but in the limit $k^+ \ll k^-$, $1/k^+ + 1/k^- \approx 1/k^+$. Note that this limit ensures that the waiting time between bursts is approximately exponentially distributed, with mean set by the only timescale left in the problem, $1/k^+$.

Let us now verify this intuition with a precise derivation to check that r/k^- is in fact the mean burst size and to obtain the full burst size distribution. Consider first a constant, known dwell time T in the active state. Transcripts are produced at a rate r per unit time, so the number of transcripts n produced during T is Poisson

distributed with mean rT , i.e.,

$$P(n | T) = \frac{(rT)^n}{n!} \exp(-rT). \quad (2.39)$$

Since the dwell time T is unobservable, we actually want $P(n)$, the dwell time distribution with no conditioning on T . Basic rules of probability theory tell us we can write $P(n)$ in terms of $P(n | T)$ as

$$P(n) = \int_0^\infty P(n | T) P(T) dT. \quad (2.40)$$

But we know the dwell time distribution $P(T)$, which is exponentially distributed according to

$$P(T) = k^- \exp(-Tk^-), \quad (2.41)$$

so $P(n)$ can be written as

$$P(n) = k^- \frac{r^n}{n!} \int_0^\infty T^n \exp[-(r + k^-)T] dT. \quad (2.42)$$

A standard integral table shows $\int_0^\infty x^n e^{-ax} dx = n!/a^{n+1}$, so

$$P(n) = \frac{k^- r^n}{(k^- + r)^{n+1}} = \frac{k^-}{k^- + r} \left(\frac{r}{k^- + r} \right)^n = \frac{k^-}{k^- + r} \left(1 - \frac{k^-}{k^- + r} \right)^n, \quad (2.43)$$

which is exactly the geometric distribution with standard parameter $\theta \equiv k^-/(k^- + r)$ and domain $n \in \{0, 1, 2, \dots\}$. The mean of the geometric distribution, with this convention, is

$$\langle n \rangle = \frac{1 - \theta}{\theta} = \left(1 - \frac{k^-}{k^- + r} \right) \frac{k^- + r}{k^-} = \frac{r}{k^-}, \quad (2.44)$$

exactly as we guessed intuitively above.

So in taking the limit $r, k^- \rightarrow \infty$, $r/k^- \equiv b$, we obtain a model which effectively has only a single promoter state, which initiates bursts at rate k^+ (transitions to the active state, in the model 4 picture). The master equation for mRNA copy number m as derived in Appendix A.1 takes the form

$$\begin{aligned} \frac{d}{dt} p(m, t) = & (m + 1) \gamma p(m + 1, t) - m \gamma p(m, t) \\ & + \sum_{m'=0}^{m-1} k_i p(m', t) \text{Geom}(m - m'; b) - \sum_{m'=m+1}^{\infty} k_i p(m, t) \text{Geom}(m' - m; b), \end{aligned} \quad (2.45)$$

where we use k_i to denote the burst initiation rate, $\text{Geom}(n; b)$ is the geometric distribution with mean b , i.e., $\text{Geom}(n; b) = \frac{1}{1+b} \left(\frac{b}{1+b}\right)^n$ (with domain over nonnegative integers as above). The first two terms are the usual mRNA degradation terms. The third term enumerates all ways the system can produce a burst of transcripts and arrive at copy number m , given that it had copy number m' before the burst. The fourth term allows the system to start with copy number m , then produce a burst and end with copy number m' . In fact this last sum has trivial m' dependence and simply enforces normalization of the geometric distribution. Carrying it out we have

$$\begin{aligned} \frac{d}{dt}p(m, t) = & (m+1)\gamma p(m+1, t) - m\gamma p(m, t) \\ & + \sum_{m'=0}^{m-1} k_i p(m', t) \text{Geom}(m-m'; b) - k_i p(m, t), \end{aligned} \quad (2.46)$$

We direct readers again to Appendix A.2 for further details. This improves on model 4 in that now the parameters are easily inferred, as we will see later, and have clean interpretations. The non-Poissonian features are attributed to the empirically well-established phenomenological picture of bursty transcription.

The big approximation in going from model 4 to 5 is that a burst is produced instantaneously rather than over a finite time. If the true burst duration is not short compared to transcription factor kinetic timescales, this could be a problem in that mean burst size in the presence and absence of repressors could change, rendering parameter inferences from the constitutive case inappropriate. Let us make some simple estimates of this.

Consider the time delay between the first and final RNAPs in a burst initiating transcription (*not* the time to complete transcripts, which potentially could be much longer.) If this timescale is short compared to the typical search timescale of transcription factors, then all is well. The estimates from deHaseth et. al. [39] put RNAP's diffusion-limited on rate around $\sim \text{few} \times 10^{-2} \text{ nM}^{-1} \text{ s}^{-1}$ and polymerase loading as high as 1 s^{-1} . Then for reasonable burst sizes of < 10 , it is reasonable to guess that bursts might finish initiating on a timescale of tens of seconds or less (with another 30-60 sec to finish elongation, but that does not matter here). A transcription factor with typical copy number of order 10 (or less) would have a diffusion-limited association rate of order $(10 \text{ sec})^{-1}$ [56]. Higher copy number TFs tend to have many binding sites over the genome, which should serve to pull them out of circulation and keep their effective association rates from rising too large. Therefore, there is *perhaps* a timescale separation possible between transcription

factor association rates and burst durations, but this assumption could very well break down, so we will have to keep it in mind when we infer repressor rates from the Jones et. al. single-cell mRNA counts data later [36].

In reflecting on these 5 models, the reader may feel that exploring a multitude of potential models just to return to a very minimal phenomenological model of bursty transcription may seem highly pedantic. But the purpose of the exercise was to examine a host of models from the literature and understand why they are insufficient, one way or another, for our purposes. Along the way we have learned that the detailed kinetics of RNAP binding and initiating transcription are probably irrelevant for setting the population distribution of mRNA. The timescales are simply too fast, and as we will see later in Figures 2.3 and 2.4, the noise seems to be governed by slower timescales. Perhaps in hindsight this is not surprising: intuitively, the degradation rate γ sets the fundamental timescale for mRNA dynamics, and any other processes that substantially modulate the mRNA distribution should not differ from γ by orders of magnitude.

2.5 Finding the “right” model: Bayesian parameter inference

In this section of the paper, we continue our program of providing one complete description of the entire broad sweep of studies that have been made in the context of the repressor-operator model, dating all the way back to the original work of Jacob and Monod and including the visionary quantitative work of Müller-Hill and collaborators [30] and up to more recent studies [33]. In addition, the aim is to reconcile the equilibrium and non-equilibrium perspectives that have been brought to bear on this problem. From Section 2.3, this reconciliation depends on a key quantitative question as codified by Eq. 2.35: does the free energy of repressor binding, as described in the equilibrium models and indirectly inferred from gene expression measurements, agree with the corresponding values of repressor binding and unbinding rates in the non-equilibrium picture, measured or inferred more directly? In this section we tackle the statistical inference problem of inferring these repressor rates from single-cell mRNA counts data. But before we can turn to the full case of simple repression, we must choose an appropriate model of the constitutive promoter and infer the parameter values in that model. This is the problem we address first.

Parameter inference for constitutive promoters

From consideration of Fano factors in the previous section, we suspect that model 5 in Figure 2.2(A), a one-state bursty model of constitutive promoters, achieves the right balance of complexity and simplicity, by allowing both Fano factor $\nu > 1$, but also by remedying, by design, the problems of parameter degeneracy that model 4 in Figure 2.2 suffered [42]. Does this stand up to closer scrutiny, namely, comparison to full mRNA distributions rather than simply their moments? We will test this thoroughly on single-cell mRNA counts for different unregulated promoters from Jones et. al. [36].

It will be instructive, however, to first consider the Poisson promoter, model 1 in Figure 2.2. As we alluded to earlier, since the Poisson distribution has a Fano factor ν strictly equal to 1, and all of the observed data in Figure 2.2(B) has Fano factor $\nu > 1$, we might already suspect that this model is incapable of fitting the data. We will verify that this is in fact the case. Using the same argument we can immediately rule out models 2 and 3 from Figure 2.2(A). These models have Fano factors $\nu \leq 1$ meaning they are underdispersed relative to the Poisson distribution. We will also not explicitly consider model 4 from Figure 2.2(A) since it was already thoroughly analyzed in [42], and since model 5 can be viewed as a special case of it.

Our objective for this section will then be to assess whether or not model 5 is quantitatively able to reproduce experimental data. In other words, if our claim is that the level of coarse graining in this model is capable of capturing the relevant features of the data, then we should be able to find values for the model parameters that can match theoretical predictions with single-molecule mRNA count distributions. A natural language for this parameter inference problem is that of Bayesian probability. We will then build a Bayesian inference pipeline to fit the model parameters to data. To gain intuition on how this analysis is done we will begin with the “wrong” model 1 in Figure 2.2(A). We will use the full dataset of single-cell mRNA counts from [36] used in Figure 2.2(B).

Model 1: Poisson promoter

For this model the master equation of interest is Eq. 2.10 with repressor set to zero, i.e.,

$$\frac{d}{dt}p_U(m)(t) = rp_U(m-1)(t) - rp_U(m)(t) + (m+1)\gamma p_U(m+1)(t) - \gamma p_U(m)(t), \quad (2.47)$$

whose steady-state solution is given by a Poisson distribution with parameter $\lambda \equiv r/\gamma$ [57]. The goal of our inference problem is then to find the probability distribution for the parameter value λ given the experimental data. By Bayes' theorem this can be written as

$$p(\lambda | D) = \frac{p(D | \lambda)p(\lambda)}{p(D)}, \quad (2.48)$$

where $D = \{m_1, m_2, \dots, m_N\}$ are the single-cell mRNA experimental counts. As is standard we will neglect the denominator $p(D)$ on the right hand side since it is independent of λ and serves only as a normalization factor.

The steady-state solution for the master equation defines the likelihood term for a single cell $p(m | \lambda)$. What this means is that for a given choice of parameter λ , under model 1 of Figure 2.2(A), we expect to observe m mRNAs in a single cell with probability

$$p(m | \lambda) = \frac{\lambda^m e^{-\lambda}}{m!}. \quad (2.49)$$

Assuming each cell's mRNA count in our dataset is independent of others, the likelihood of the full inference problem $p(D | \lambda)$ is simply a product of the single cell likelihoods given by Eq. 2.49 above, so

$$p(D | \lambda) = \prod_{k=1}^N \frac{\lambda^{m_k} e^{-\lambda}}{m_k!}. \quad (2.50)$$

To proceed we need to specify a prior distribution $p(\lambda)$. In this case we are extremely data-rich, as the dataset from Jones et. al [36] has of order 1000-3000 single-cell measurements for each promoter, so our choice of prior matters little here, as long as it is sufficiently broad. For details on the prior selection we refer the reader to Appendix A.3. For our purpose here it suffices to specify that we use as prior a Gamma distribution. This particular choice of prior introduces two new parameters, α and β , which parametrize the gamma distribution itself, which we use to encode the range of λ values we view as reasonable. Recall λ is the mean steady-state mRNA count per cell, which *a priori* could plausibly be anywhere from 0 to a few hundred. $\alpha = 1$ and $\beta = 1/50$ achieve this, since the gamma distribution is strictly positive with mean α/β and standard deviation $\sqrt{\alpha}/\beta$.

As detailed in Appendix A.3 this particular choice of prior is known as the *conjugate* prior for a Poisson likelihood. Conjugate priors have the convenient properties that a closed form exists for the posterior distribution $p(\lambda | D)$ - unusual in Bayesian inference problems - and the closed form posterior takes the same form as the prior.

For our case of a Poisson distribution likelihood with its Gamma distribution conjugate prior, the posterior distribution is also a Gamma distribution [52]. Specifically the two parameters α' and β' for this posterior distribution take the form $\alpha' = \alpha + \bar{m}N$ and $\beta' = \beta + N$, where we defined the sample mean $\bar{m} = \frac{1}{N} \sum_{k=1}^N m_k$ for notational convenience, and N is the number of cells in our dataset. Furthermore, given that N is $O(10^3)$ and $\langle m \rangle \gtrsim 0.1$ for all promoters measured in [36] our data easily overwhelms the choice of prior, and allows us to approximate the Gamma distribution with a Gaussian distribution with mean \bar{m} and variance \bar{m}/N with marginal errors. As an example with real numbers, for the *lacUV5* promoter, Jones et. al [36] measured 2648 cells with an average mRNA count per cell of $\bar{m} \approx 18.7$. For this case our posterior distribution $P(\lambda \mid D)$ would be a Gaussian distribution with mean $\mu = 18.7$, and a standard deviation $\sigma \approx 0.08$. This suggests we have inferred our model's one parameter to a precision of order 1%.

We remind the reader that we began this section claiming that the Poisson model was “wrong” since it could not reproduce features of the data such as a Fano factor > 1 . The fact that we obtain such a narrow posterior distribution for our parameter $P(\lambda \mid D)$ does not equate to the model being adequate to describe the data. What this means is that given the data D , only values in a narrow range are remotely plausible for the parameter λ , but a narrow posterior distribution does not necessarily mean the model accurately depicts reality. As we will see later in Figure 2.3 after exploring the bursty promoter model, indeed the correspondence when contrasting the Poisson model with the experimental data is quite poor.

Model 5 - Bursty promoter

Let us now consider the problem of parameter inference for model five from Figure 2.1(C). As derived in Appendix A.2, the steady-state mRNA distribution in this model is a negative binomial distribution, given by

$$p(m) = \frac{\Gamma(m + k_i)}{\Gamma(m + 1)\Gamma(k_i)} \left(\frac{1}{1 + b} \right)^{k_i} \left(\frac{b}{1 + b} \right)^m, \quad (2.51)$$

where b is the mean burst size and k_i is the burst rate in units of the mRNA degradation rate γ . As sketched earlier, to think of the negative binomial distribution in terms of an intuitive “story,” in the precise meaning of [61], we imagine the arrival of bursts as a Poisson process with rate k_i , where each burst has a geometrically-distributed size with mean size b .

As for the Poisson promoter model, this expression for the steady-state mRNA distribution is exactly the likelihood we want to use when stating Bayes theorem. Again denoting the single-cell mRNA count data as $D = \{m_1, m_2, \dots, m_N\}$, here Bayes' theorem takes the form

$$p(k_i, b \mid D) \propto p(D \mid k_i, b)p(k_i, b). \quad (2.52)$$

We already have our likelihood – the product of N negative binomials as Eq. 2.51 – so we only need to choose priors on k_i and b . For the datasets from [36] that we are analyzing, as for the Poisson promoter model above we are still data-rich so the prior's influence remains weak, but not nearly as weak because the dimensionality of our model has increased from one parameter to two. Details on the arguments behind our prior distribution selection are left for Appendix A.3. We state here that the natural scale to explore these parameters is logarithmic. This is commonly the case for parameters for which our previous knowledge based on our domain expertise spans several orders of magnitude. For this we chose log-normal distributions for both k_i and b . Details on the mean and variance of these distributions can be found in Appendix A.3.

We carried out Markov-Chain Monte Carlo (MCMC) sampling on the posterior of this model, starting with the constitutive *lacUV5* dataset from [36]. The resulting MCMC samples are shown in Figure 2.3(A). In contrast to the active/inactive constitutive model considered in [42] (nonequilibrium model 4 in Figure 2.2(A)), this model is well-identified with both parameters determined to a fractional uncertainty of 5-10%. The strong correlation reflects the fact that their product sets the mean of the mRNA distribution, which is tightly constrained by the data, but there is weak “sloppiness” [62] along a set of values with a similar product.

Having found the model's posterior to be well-identified as with the Poisson promoter, the next step is to compare both models with experimental data. To do this for the case of the bursty promoter, for each of the parameter samples shown in Figure 2.3(A) we generated negative binomial-distributed mRNA counts. As MCMC samples parameter space proportionally to the posterior distribution, this set of random samples span the range of possible values that we would expect given the correspondence between our theoretical model and the experimental data. A similar procedure can be applied to the Poisson promoter. To compare so many samples with the actual observed data, we can use empirical cumulative distribution functions (ECDF) of the distribution quantiles. This representation is shown in

Figure 2.3(B). In this example, the median for each possible mRNA count for the Poisson distribution is shown as a dark green line, while the lighter green contains 95% of the randomly generated samples. This way of representing the fit of the model to the data gives us a sense of the range of data we might consider plausible, under the assumption that the model is true. For this case, as we expected given our premise of the Poisson promoter being wrong, it is quite obvious that the observed data, plotted in black is not consistent with the Poisson promoter model. An equivalent plot for the bursty promoter model is shown in blue. Again the darker tone shows the median, while the lighter color encompasses 95% of the randomly generated samples. Unlike the Poisson promoter model, the experimental ECDF closely tracks the posterior predictive ECDF, indicating this model is actually able to generate the observed data and increasing our confidence that this model is sufficient to parametrize the physical reality of the system.

The commonly used promoter sequence *lacUV5* is our primary target here, since it forms the core of all the simple repression constructs of [36] that we consider in Section 2.5. Nevertheless, we thought it wise to apply our bursty promoter model to the other 17 unregulated promoters available in the single-cell mRNA count dataset from [36] as a test that the model is capturing the essential phenomenology. If the model fit well to all the different promoters, this would increase our confidence that it would serve well as a foundation for inferring repressor kinetics later in Section 2.5. Conversely, were the model to fail on more than a couple of the other promoters, it would give us pause.

Figure 2.3(C) shows the results, plotting the posterior distribution from individually MCMC sampling all 18 constitutive promoter datasets from [36]. To aid visualization, rather than plotting samples for each promoter’s posterior as in Figure 2.3(A), for each posterior we find and plot the curve that surrounds the 95% highest probability density region. What this means is that each contour encloses approximately 95% of the samples, and thus 95% of the probability mass, of its posterior distribution. Theory-experiment comparisons, shown in Figure A.5 in Appendix A.3, display a similar level of agreement between data and predictive samples as for the bursty model with *lacUV5* in Figure 2.3(B).

One interesting feature from Figure 2.3(C) is that burst rate varies far more widely, over a range of $\sim 10^2$, than burst size, confined to a range of $\lesssim 10^1$ (and with the exception of promoter 6, just a span of 3 to 5-fold). This suggests that k_i , not b , is the key dynamic variable that promoter sequence tunes.

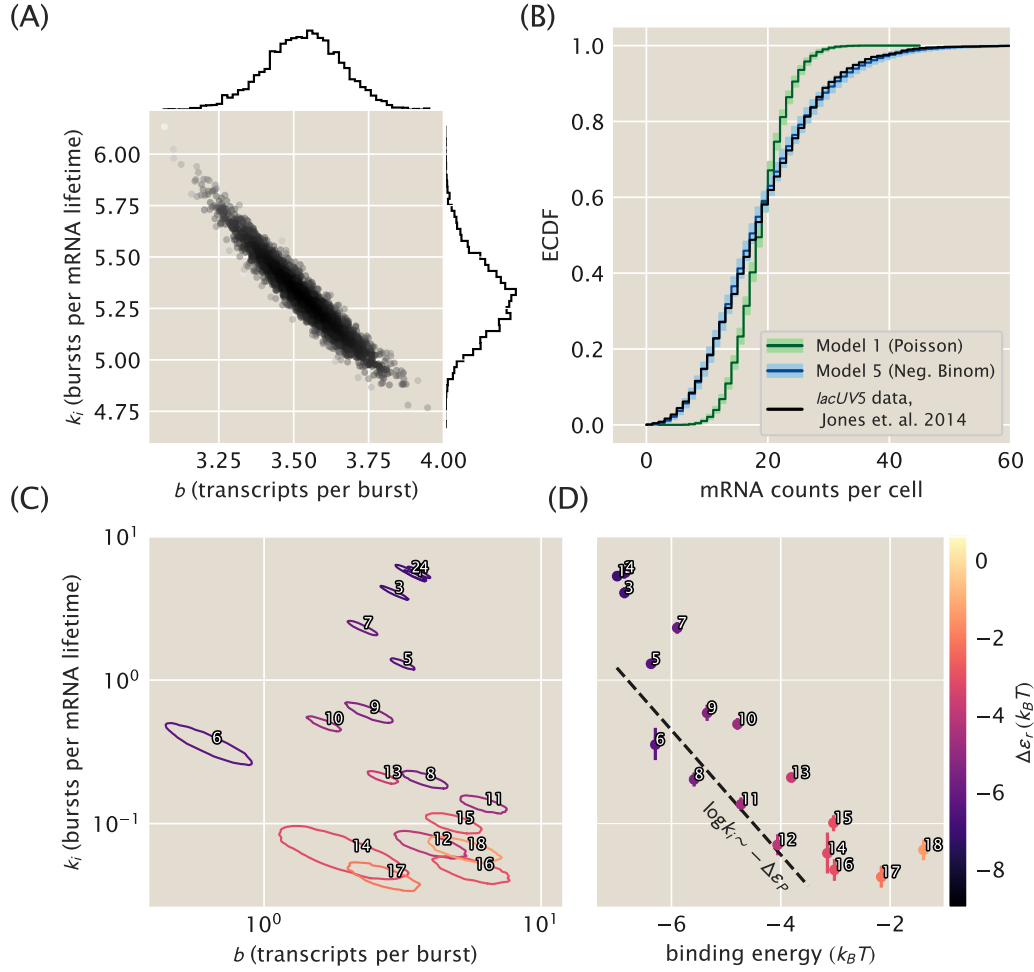


Figure 2.3: Constitutive promoter posterior inference and model comparison.

(A) The joint posterior density of model 5, the bursty promoter with negative binomially-distributed steady state, is plotted with MCMC samples. 1D marginal probability densities are plotted as flanking histograms. The model was fit on *lacUV5* data from [36]. (B) The empirical distribution function (ECDF) of the observed population distribution of mRNA transcripts under the control of a constitutive *lacUV5* promoter is shown in black. The median posterior predictive ECDFs for models (1), Poisson, and (5), negative binomial, are plotted in dark green and dark blue, respectively. Lighter green and blue regions enclose 95% of all posterior predictive samples from their respective models. Model (1) is in obvious contradiction with the data while model (5) is not. Single-cell mRNA count data is again from [36]. (C) Joint posterior distributions for burst rate k_i and mean burst size b for 18 unregulated promoters from [36]. Each contour indicates the 95% highest posterior probability density region for a particular promoter. Note that the vertical axis is shared with (D). (D) Plots of the burst rate k_i vs. the binding energy for each promoter as predicted in [54]. The dotted line shows the predicted slope according to Eq. 2.55, described in text. Each individual promoter is labeled with a unique number in both (C) and (D) for cross comparison and for comparison with Figure 2.2(B).

Connecting inferred parameters to prior work

It is interesting to connect these inferences on k_i and b to the work of [54], where these same 18 promoters were considered through the lens of the three-state equilibrium model (model 2 in Figure 2.1(B)) and binding energies $\Delta\epsilon_P$ were predicted from an energy matrix model derived from [53]. As previously discussed the thermodynamic models of gene regulation can only make statements about the mean gene expression. This implies that we can draw the connection between both frameworks by equating the mean mRNA $\langle m \rangle$. This results in

$$\langle m \rangle = \frac{k_i b}{\gamma} = \frac{r}{\gamma} \frac{\frac{P}{N_{NS}} \exp(-\beta\Delta\epsilon_P)}{1 + \frac{P}{N_{NS}} \exp(-\beta\Delta\epsilon_P)}. \quad (2.53)$$

By taking the weak promoter approximation for the equilibrium model ($P/N_{NS} \exp(-\beta\Delta\epsilon_P) \ll 1$) results in

$$\langle m \rangle = \frac{k_i b}{\gamma} = \frac{r}{\gamma} \frac{P}{N_{NS}} \exp(-\beta\Delta\epsilon_P), \quad (2.54)$$

valid for all the binding energies considered here.

Given this result, how are the two coarse-grainings related? A quick estimate can shed some light. Consider for instance the *lacUV5* promoter, which we see from Figure 2.3(A) has $k_i/\gamma \sim b \sim \text{few}$, from Figure 2.3(B) has $\langle m \rangle \sim 20$, and from [54] has $\beta\Delta\epsilon_P \sim -6.5$. Further we generally assume $P/N_{NS} \sim 10^{-3}$ since $N_{NS} \approx 4.6 \times 10^6$ and $P \sim 10^3$. After some guess-and-check with these values, one finds the only association that makes dimensional sense and produces the correct order-of-magnitude for the known parameters is to take

$$\frac{k_i}{\gamma} = \frac{P}{N_{NS}} \exp(-\beta\Delta\epsilon_P) \quad (2.55)$$

and

$$b = \frac{r}{\gamma}. \quad (2.56)$$

Figure 2.3(D) shows that this linear scaling between $\ln k_i$ and $-\beta\Delta\epsilon_P$ is approximately true for all 18 constitutive promoters considered. The plotted line is simply Eq. 2.55 and assumes $P \approx 5000$.

While the associations represented by Eq. 2.55 and Eq. 2.56 appear to be borne out by the data in Figure 2.3, we do not find the association of parameters they imply to be intuitive. We are also cautious to ascribe too much physical reality to the parameters. Indeed, part of our point in comparing the various constitutive

promoter models is to demonstrate that these models each provide an internally self-consistent framework that adequately describes the data, but attempting to translate between models reveals the dubious physical interpretation of their parameters.

We mention one further comparison, between our inferred parameters and the work of Chong et. al. [47], which is interesting and puzzling. Beautiful experiments in [47] convincingly argue that supercoiling accumulated from the production of mRNA transcripts is key in setting the burstiness of mRNA production. In their model, this supercoiling occurs on the scale of ~ 100 kb domains of DNA. This suggests that all genes on a given domain should burst in synchrony, and that the difference between highly and lowly expressed genes is the *size* of transcriptional bursts, not the *time between* bursts. But here, all burst sizes we infer in Figure 2.3(C) are comparable and burst rates vary wildly. It is not immediately clear how to square this circle. Furthermore, Figure 7E in [47] reports values of the quantity they label β/α and we label k^+/k^- in model 4 from Figure 2.2. In contrast to the findings of [42], Chong et. al. do not find $k^+/k^- \ll 1$ for most of the genes they consider. This begs the question: is the *galK* chromosomal locus used for the reporter constructs in [42] and [36] merely an outlier, or is there a deeper puzzle here waiting to be resolved? Without more apples-to-apples data we can only speculate, and we leave it as an intriguing open question for the field.

Despite such puzzles, our goal here is not to unravel the mysterious origins of burstiness in transcription. Our remaining task in this work is a determination of the physical reality of equilibrium binding energies in Figure 2.1, as codified by the equilibrium-nonequilibrium equivalence of Eq. 2.35. For our phenomenological needs here model 5 in Figure 2.2 is more than adequate: the posterior distributions in Figure 2.3(C) are cleanly identifiable and the predictive checks in Figure A.5 indicate no discrepancies between the model and the mRNA single-molecule count data of [36]. Of the models we have considered it is unique in satisfying both these requirements. So we will happily use it as a foundation to build upon in the next section when we add regulation.

Transcription factor kinetics can be inferred from single-cell mRNA distribution measurements

Building the model and performing parameter inference

Now that we have a satisfactory model in hand for constitutive promoters, we would like to return to the main thread: can we reconcile the equilibrium and

nonequilibrium models by putting to the test Eq. 2.35, the correspondence between indirectly inferred equilibrium binding energies and nonequilibrium kinetic rates? To make this comparison, is it possible to infer repressor binding and unbinding rates from mRNA distributions over a population of cells as measured by single-molecule Fluorescence *in situ* Hybridization in [36]? If so, how do these inferred rates compare to direct single-molecule measurements such as from [56] and to binding energies such as from [33] and [32], which were inferred under the assumptions of the equilibrium models in Figure 2.1(B)? And can this comparison shed light on the unreasonable effectiveness of the equilibrium models, for instance, in their application in [35, 63]?

As we found in Section 2.4, for our purposes the “right” model of a constitutive promoter is the bursty picture, model five in Figure 2.2(A). Therefore our starting point here is the analogous model with repressor added, model 5 in Figure 2.1(C). For a given repressor binding site and copy number, this model has four rate parameters to be inferred: the repressor binding and unbinding rates k_R^+ , and k_R^- , the initiation rate of bursts, k_i , and the mean burst size b (we nondimensionalize all of these by the mRNA degradation rate γ).

Before showing the mathematical formulation of our statistical inference model, we would like to sketch the intuitive structure. The dataset from [36] we consider consists of single-cell mRNA counts data of nine different conditions, spanning several combinations of three unique repressor binding sites and four unique repressor copy numbers. We assume that the values of k_i and b are known, since we have already cleanly inferred them from constitutive promoter data, and further we assume that these values are the same across datasets with different repressor binding sites and copy numbers. In other words, we assume that the regulation of the transcription factor does not affect the mean burst size nor the burst initiation rate. The regulation occurs as the promoter is taken away from the transcriptionally active state when the promoter is bound by repressor. We assume that there is one unbinding rate parameter for each repressor binding site, and likewise one binding rate for each unique repressor copy number. This makes our model seven dimensional, or nine if one counts k_i and b as well. Note that we use only a subset of the datasets from Jones et. al. [36], as discussed more in Appendix A.3.

Formally now, denote the set of seven repressor rates to be inferred as

$$\vec{k} = \{k_{Oid}^-, k_{O1}^-, k_{O2}^-, k_{0.5}^+, k_1^+, k_2^+, k_{10}^+\}, \quad (2.57)$$

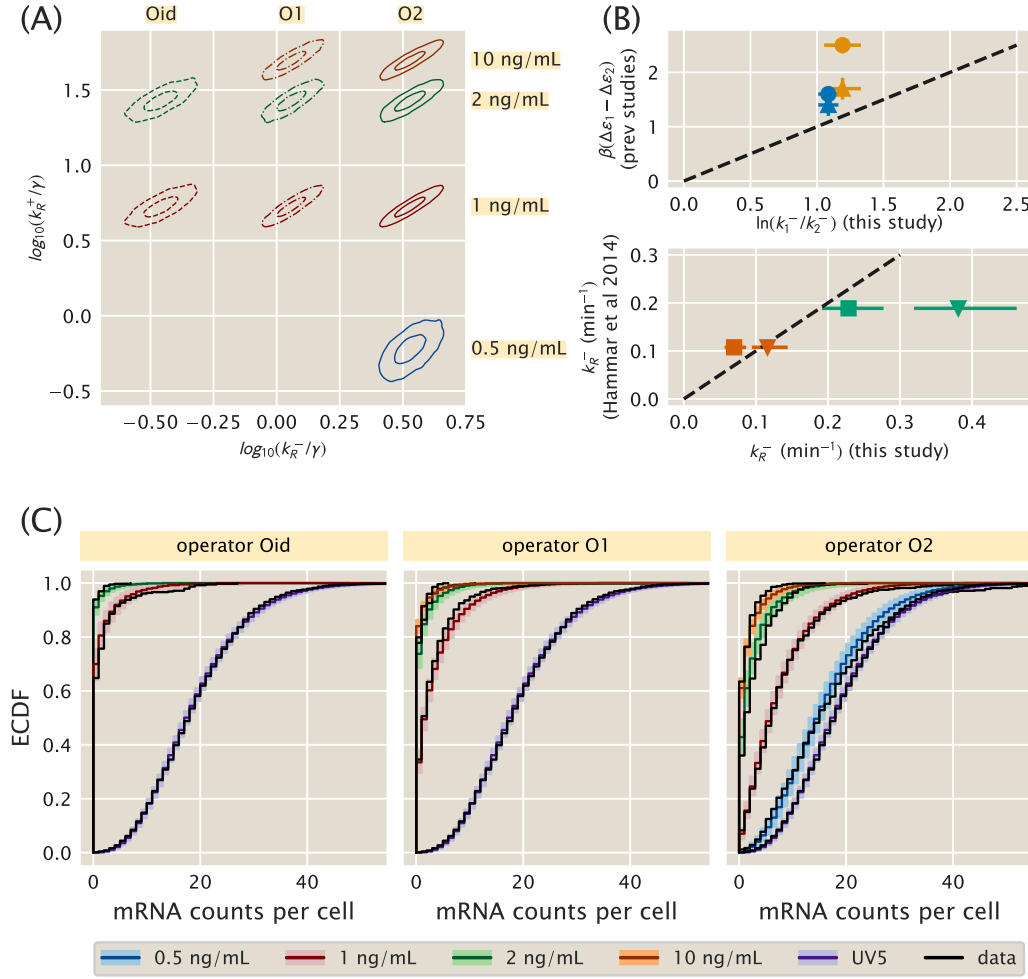


Figure 2.4: Simple repression parameter inference and comparison. (A) Contours which enclose 50% and 95% of the posterior probability mass are shown for each of several 2D slices of the 9D posterior distribution. The model assumes one unbinding rate for each operator (Oid, O1, O2) and one binding rate for each aTc induction concentration (corresponding to an unknown mean repressor copy number). (B, upper) Ratios of our inferred unbinding rates are compared with operator binding energy differences measured by Garcia and Phillips [33] (triangles) and Razo-Mejia et. al. [32] (circles). Blue glyphs compare O2-O1, while orange compare O1-Oid. Points with perfect agreement would lie on the dotted line. (B, lower) Unbinding rates for O1 (cyan) and Oid (red) inferred in this work are compared with single-molecule measurements from Hammar et. al. [56]. We plot the comparison assuming illustrative mRNA lifetimes of $\gamma^{-1} = 3$ min (triangles) or $\gamma^{-1} = 5$ min (squares). Dotted line is as in upper panel. (C) Theory-experiment comparison are shown for each of the datasets used in the inference of the model in (A). Observed single-molecule mRNA counts data from [36] are plotted as black lines. The median of the randomly generated samples for each condition is plotted as a dark colored line. Lighter colored bands enclose 95% of all samples for a given operator/repressor copy number pair. The unregulated promoter, *lacUV5*, is shown with each as a reference.

where subscripts for dissociation rates k^- indicate the different repressor binding sites, and subscripts for association rates k^+ indicate the concentration of the small-molecule that controlled the expression of the LacI repressor (see Appendix A.3). This is because for this particular dataset the repressor copy numbers were not measured directly, but it is safe to assume that a given concentration of the inducer resulted in a specific mean repressor copy number [63]. Also note that the authors of [36] report estimates of LacI copy number per cell rather than direct measurements. However, these estimates were made assuming the validity of the equilibrium models in Figure 2.1, and since testing these models is our present goal, it would be circular logic if we were to make the same assumption. Therefore we will make no assumptions about the LacI copy number for a given inducer concentrations.

Having stated the problem, Bayes' theorem reads

$$p(\vec{k}, k_i, b \mid D) \propto p(D \mid \vec{k}, k_i, b) p(\vec{k}, k_i, b), \quad (2.58)$$

where D is again the set of all N observed single-cell mRNA counts across the various conditions. We assume that individual single-cell measurements are independent so that the likelihood factorizes as

$$p(D \mid \vec{k}, k_i, b) = \prod_{j=1}^N p(m \mid \vec{k}, k_i, b) = \prod_{j=1}^N p(m \mid k_j^+, k_j^-, k_i, b) \quad (2.59)$$

where k_j^\pm represent the appropriate binding and unbinding rates out of \vec{k} for the j -th measured cell. The probability $p(m \mid k_j^+, k_j^-, k_i, b)$ appearing in the last expression is exactly Eq. A.184, the steady-state distribution for our bursty model with repression derived in Section A.2, which for completeness we reproduce here as

$$p(m \mid k_R^+, k_R^-, k_i, b) = \frac{\Gamma(\alpha + m) \Gamma(\beta + m) \Gamma(k_R^+ + k_R^-) b^m}{\Gamma(\alpha) \Gamma(\beta) \Gamma(k_R^+ + k_R^- + m) m!} \quad (2.60)$$

$$\times {}_2F_1(\alpha + m, \beta + m, k_R^+ + k_R^- + m; -b).$$

where ${}_2F_1$ is the confluent hypergeometric function of the second kind and α and β , defined for notational convenience, are

$$\alpha = \frac{1}{2} \left(k_i + k_R^- + k_R^+ + \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right) \quad (2.61)$$

$$\beta = \frac{1}{2} \left(k_i + k_R^- + k_R^+ - \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right).$$

This likelihood is rather inscrutable. We did not find any of the known analytical approximations for ${}_2F_1$ terribly useful in gaining intuition, so we instead resorted to numerics. One insight we found was that for very strong or very weak repression, the distribution in Eq. 2.60 is well approximated by a negative binomial with burst size b and burst rate k_i equal to their constitutive *lacUV5* values, except with k_i multiplied by the fold-change $(1 + k_R^+/k_R^-)^{-1}$. In other words, once again only the ratio k_R^+/k_R^- was detectable. But for intermediate repression, the distribution was visibly broadened with Fano factor greater than $1+b$, the value for the corresponding constitutive case. This indicates that the repressor rates had left an imprint on the distribution, and perhaps intuitively, this intermediate regime occurs for values of k_R^\pm comparable to the burst rate k_i . Put another way, if the repressor rates are much faster or much slower than k_i , then there is a timescale separation and effectively only one timescale remains, $k_i (1 + k_R^+/k_R^-)^{-1}$. Only when all three rates in the problem are comparable does the mRNA distribution retain detectable information about them.

Next we specify priors. As for the constitutive model, weakly informative log-normal priors are a natural choice for all our rates. We found that if the priors were too weak, our MCMC sampler would often become stuck in regions of parameter space with very low probability density, unable to move. We struck a balance in choosing our prior widths between helping the sampler run while simultaneously verifying that the marginal posteriors for each parameter were not artificially constrained or distorted by the presence of the prior. All details for our prior distributions are listed in Appendix A.3.

We ran MCMC sampling on the full nine dimensional posterior specified by this model. To attempt to visualize this object, in Figure 2.4(A) we plot several two-dimensional slices as contour plots, analogous to Figure 2.3(C). Each of these nine slices corresponds to the (k_R^+, k_R^-) pair of rates for one of the conditions from the dataset used to fit the model and gives a sense of the uncertainty and correlations in the posterior. We note that the 95% uncertainties of all the rates span about ~ 0.3 log units, or about a factor of two, with the exception of $k_{0.5}^+$, the association rate for the lowest repressor copy number which is somewhat larger.

Comparison with prior measurements of repressor binding energies

Our primary goal in this work is to reconcile the kinetic and equilibrium pictures of simple repression. Towards this end we would like to compare the repressor

kinetic rates we have inferred with the repressor binding energies inferred through multiple methods in [33] and [32]. If the agreement is close, then it suggests that the equilibrium models are not wrong and the repressor binding energies they contain correspond to physically real free energies, not mere fit parameters.

Figure 2.4(B) shows both comparisons, with the top panel comparing to equilibrium binding energies and the bottom panel comparing to single-molecule measurements. First consider the top panel and its comparison between repressor kinetic rates and binding energies. As described in section 2.3, if the equilibrium binding energies from [33] and [32] indeed are the physically real binding energies we believe them to be, then they should be related to the repressor kinetic rates via Eq. 2.35, which we restate here,

$$\Delta F_R = \beta \Delta \varepsilon_R - \log(R/N_{NS}) = -\log(k_R^+/k_R^-). \quad (2.62)$$

Assuming mass action kinetics implies that k_R^+ is proportional to repressor copy number R , or more precisely, it can be thought of as repressor copy number times some intrinsic per molecule association rate. But since R is not directly known for our data from [36], we cannot use this equation directly. Instead we can consider two different repressor binding sites and compute the *difference* in binding energy between them, since this difference depends only on the unbinding rates and not on the binding rates. This can be seen by evaluating Eq. 2.62 for two different repressor binding sites, labeled (1) and (2), but with the same repressor copy number R , and taking the difference to find

$$\Delta F_R^{(1)} - \Delta F_R^{(2)} = \beta \Delta \varepsilon_1 - \beta \Delta \varepsilon_2 = -\log(k_R^+/k_1^-) + \log(k_R^+/k_2^-), \quad (2.63)$$

or simply

$$\beta \Delta \varepsilon_1 - \beta \Delta \varepsilon_2 = \log(k_2^-/k_1^-). \quad (2.64)$$

The left and right hand sides of this equation are exactly the horizontal and vertical axes of the top panel of Figure 2.4. Since we inferred rates for three repressor binding sites (O1, O2, and Oid), there are only two independent differences that can be constructed, and we arbitrarily chose to plot O2-O1 and O1-Oid in Figure 2.4(B). Numerically, we compute values of k_{O1}^-/k_{Oid}^- and k_{O2}^-/k_{O1}^- directly from our full posterior samples, which conveniently provides uncertainties as well, as detailed in Appendix A.3. We then compare these log ratios of rates to the binding energy differences $\Delta \varepsilon_{O1} - \Delta \varepsilon_{Oid}$ and from $\Delta \varepsilon_{O2} - \Delta \varepsilon_{O1}$ as computed from the values from both [33] and [32]. Three of the four values are within $\sim 0.5 k_B T$ of the diagonal

representing perfect agreement, which is comparable to the $\sim \text{few} \times 0.1 k_B T$ variability between the independent determinations of the same quantities between [33] and [32]. The only outlier involves Oid measurements from [32], and as the authors of [32] note, this is a difficult measurement of low fluorescence signal against high background since Oid represses so strongly. We are therefore inclined to regard the failure of this point to fall near the diagonal as a testament to the difficulty of the measurement and not as a failure of our theory.

On the whole then, we regard this as striking confirmation of the validity of the equilibrium models. Their lynchpin parameter is a phenomenological free energy of repressor binding that has previously only been inferred indirectly. Our result shows that the microscopic interpretation of this free energy, as the log of a ratio of transition rates, does indeed hold true to within the inherent uncertainties that remain in the entire theory-experiment dialogue.

Comparison with prior measurements of repressor kinetics

In the previous section we established the equivalence between the equilibrium models' binding energies and the repressor kinetics we infer from mRNA population distributions. But one might worry that the repressor rates we infer from mRNA distributions are *themselves* merely fitting parameters and that they do not actually correspond to the binding and unbinding rates of the repressor in vivo. To verify that this is not the case, we next compare our kinetic rates with a different measurement of the same rates using a radically different method: single molecule measurements as performed in Hammar et. al. [56]. This is plotted in the lower panel of Figure 2.4(B).

Since we do not have access to repressor copy number for either the single-cell mRNA data from [36] or the single-molecule data from [56], we cannot make an apples-to-apples comparison of association rates k_R^+ . Further, while Hammar et. al. directly measure the dissociation rates k_R^- , our inference procedure returns k_R^-/γ , i.e., the repressor dissociation rate nondimensionalized by the mRNA degradation rate γ . So to make the comparison, we must make an assumption for the value of γ since it was not directly measured. For most mRNAs in *E. coli*, quoted values for the typical mRNA lifetime γ^{-1} range between about 2.5 min [64] to 8 min. We chose $\gamma^{-1} = 3$ min and $\gamma^{-1} = 5$ min as representative values and plot a comparison of k_{O1}^- and k_{Oid}^- from our inference with corresponding values reported in [56] for both these choices of γ .

The degree of quantitative agreement in the lower panel of Figure 2.4(B) clearly

depends on the precise choice of γ . Nevertheless we find this comparison very satisfying, when two wildly different approaches to a measurement of the same quantity yield broadly compatible results. We emphasize the agreement between our rates and the rates reported in [56] for any reasonable γ : values differ by at most a factor of 2 and possibly agree to within our uncertainties of 10-20%. From this we feel confident asserting that the parameters we have inferred from Jones et. al.'s single-cell mRNA counts data do in fact correspond to repressor binding and unbinding rates, and therefore our conclusions on the agreement of these rates with binding energies from [33] and [32] are valid.

Model checking

In Figure 2.3(B) we saw that the simple Poisson model of a constitutive promoter, despite having a well behaved posterior, was clearly insufficient to describe the data. It behooves us to carry out a similar check for our model of simple repression, codified by Eq. 2.60 for the steady-state mRNA copy number distribution. As derived in Sections 2.3 and 2.4, we have compelling theoretical reasons to believe it is a good model, but if it nevertheless turned out to be badly contradicted by the data we should like to know.

The details are deferred to Appendix A.3, and here we only attempt to summarize the intuitive ideas, as detailed at greater length by Jaynes [65] as well as Gelman and coauthors [52, 66]. From our samples of the posterior distribution, plotted in Figure 2.4(A), we generate many replicate data using a random number generator. In Figure 2.4(C), we plot empirical cumulative distribution functions of the middle 95% quantiles of these replicate data with the actual experimental data from Jones et. al. [36] overlaid, covering all ten experimental conditions spanning repressor binding sites and copy numbers (as well as the constitutive baseline UV5).

The purpose of Figure 2.4(C) is simply a graphical, qualitative assessment of the model: do the experimental data systematically disagree with the simulated data, which would suggest that our model is missing important features? A further question is not just whether there is a detectable difference between simulated and experimental data, but whether this difference is likely to materially affect the conclusions we draw from the posterior in Figure 2.4(A). More rigorous and quantitative statistical tests are possible [52], but their quantitiveness does not necessarily make them more useful. As stated in [66], we often find this graphical comparison more enlightening because it better engages our intuition for the model,

not merely telling *if* the model is wrong but suggesting *how* the model may be incomplete.

Our broad brush takeaway from Figure 2.4(C) is overall of good agreement. There are some oddities, in particular the long tails in the data for Oid, 1 ng/mL, and O2, 0.5 ng/mL. The latter is especially odd since it extends beyond the tail of the unregulated UV5 distribution. This is a relatively small number of cells, however, so whether this is a peculiarity of the experimental data, a statistical fluke of small numbers, or a real biological effect is unclear. It is conceivable that there is some very slow timescale switching dynamics that could cause this bimodality, although it is unclear why it would only appear for specific repressor copy numbers. There is also a small offset between experiment and simulation for O2 at the higher repressor copy numbers, especially at 2 and 10 ng/mL. From the estimate of repressor copy numbers from [36], it is possible that the repressor copy numbers here are becoming large enough to partially invalidate our assumption of a separation of timescales between burst duration and repressor association rate. Another possibility is that the very large number of zero mRNA counts for Oid, 2 ng/mL is skewing its partner datasets through the shared association rate. None of these fairly minor quibbles cause us to seriously doubt the overall correctness of our model, which further validates its use to compare the equilibrium models' binding energies to the nonequilibrium models' repressor kinetics, as we originally set out to do.

2.6 Discussion and future work

The study of gene expression is one of the dominant themes of modern biology, made all the more urgent by the dizzying pace at which genomes are being sequenced. But there is a troubling Achilles heel buried in all of that genomic data, which is our inability to find and interpret regulatory sequence. In many cases, this is not possible even qualitatively, let alone the possibility of quantitative dissection of the regulatory parts of genomes in a predictive fashion. Other recent work has tackled the challenge of finding and annotating the regulatory part of genomes [5, 28]. Once we have determined the architecture of the regulatory part of the genome, we are then faced with the next class of questions which are sharpened by formulating them in mathematical terms, namely, what are the input-output properties of these regulatory circuits and what knobs control them?

The present work has tackled that question in the context of the first regulatory architecture hypothesized in the molecular biology era, namely, the repressor-operator

model of Jacob and Monod [1]. Regulation in that architecture is the result of a competition between a repressor which inhibits transcription and RNAP polymerase which undertakes it. Through the labors of generations of geneticists, molecular biologists and biochemists, an overwhelming amount of information and insight has been garnered into this simple regulatory motif, licensing it as what one might call the “hydrogen atom” of regulatory biology. It is from that perspective that the present paper explores the extent to which some of the different models that have been articulated to describe that motif allow us to understand both the average level of gene expression found in a population of cells, the intrinsic cell-to-cell variability, and the full gene expression distribution found in such a population as would be reported in a single molecule mRNA Fluorescence *in situ* Hybridization experiment, for example.

Our key insights can be summarized as follows. First, as shown in Figure 2.1, the mean expression in the simple repression architecture is captured by a master curve in which the action of repressor and the details of the RNAP interaction with the promoter appear separately and additively in an effective free energy. Interestingly, as has been shown elsewhere in the context of the Monod-Wyman-Changeux model, these kinds of coarse-graining results are an exact mathematical result and do not constitute hopeful approximations or biological naivete [32, 35]. To further dissect the relative merits of the different models, we must appeal to higher moments of the gene expression probability distribution. To that end, our second set of insights focus on gene expression noise, where it is seen that a treatment of the constitutive promoter already reveals that some models have Fano factors (variance/mean) that are less than one, at odds with any and all experimental data that we are aware of [36, 59]. This theoretical result allows us to directly discard a subset of the models (models 1-3 in Figure 2.2(A)) since they cannot be reconciled with experimental observations. The two remaining models (models 4 and 5 in Figure 2.2) appear to contain enough microscopic realism to be able to reproduce the data. A previous exploration of model 4 demonstrated the “sloppy” [62] nature of the model in which data on single-cell mRNA counts alone cannot constrain the value of all parameters simultaneously [42]. Here we demonstrate that the proposed one-state bursty promoter model (model 5 in Figure 2.2(A)) emerges as a limit of the commonly used two-state promoter model [19, 23, 36, 57, 59]. We put the idea to the test that this level of coarse-graining is rich enough to reproduce previous experimental observations. In particular we perform Bayesian inference to determine the two parameters describing the full steady-state mRNA distribution,

finding that the model is able to provide a quantitative description of a plethora of promoter sequences with different mean levels of expression and noise.

With the results of the constitutive promoter in hand, we then fix the parameters associated with this class of promoters and use them as input for evaluating the noise in gene expression for the simple repression motif itself. This allows us to provide a single overarching analysis of both the constitutive and simple repression architectures using one simple model and corresponding set of self-consistent parameters, demonstrating not only a predictive framework, but also reconciling the equilibrium and non-equilibrium views of the same simple repression constructs. More specifically, we obtained values for the transcription factor association and dissociation rates by performing Bayesian inference on the full mRNA distribution for data obtained from simple-repression promoters with varying number of transcription factors per cell and affinity of such transcription factors for the binding site. The free energy value obtained from these kinetic rates – computed as the log ratio of the rates – agrees with previous inferences performed only from mean gene expression measurements, that assumed an equilibrium rather than a kinetic framework [32, 33].

It is interesting to speculate what microscopic details are being coarse-grained by our burst rate and burst size in Figure 2.2, model 5. Chromosomal locus is one possible influence we have not addressed in this work, as all the single-molecule mRNA data from [36] that we considered was from a construct integrated at the *galK* locus. The results of [47] indicate that transcription-induced supercoiling contributes substantially in driving transcriptional bursting, and furthermore, their Figure 7 suggests that the parameters describing the rate, duration, and size of bursts vary substantially for transcription from different genomic loci. Although the authors of [67] do not address noise, they note enormous differences in mean expression levels when an identical construct is integrated at different genomic loci. The authors of [68] attribute noise and burstiness in their single-molecule mRNA data to the influence of different sigma factors, which is a reasonable conclusion from their data. Could the difference also be due to the different chromosomal locations of the two operons? What features of different loci are and are not important? Could our preferred coarse-grained model capture the variability across different loci? If so, and we were to repeat the parameter inference as done in this work, is there a simple theoretical model we could build to understand the resulting parameters?

In summary, this work took up the challenge of exploring the extent to which a

single specific mechanistic model of the simple-repression regulatory architecture suffices to explain the broad sweep of experimental data for this system. Pioneering early experimental efforts from the Müller-Hill lab established the simple-repression motif as an arena for the quantitative dissection of regulatory response in bacteria, with similar beautiful work emerging in examples such as the *ara* and *gal* operons as well [29, 30, 69–73]. In light of a new generation of precision measurements on these systems, the definition of what it means to understand them can now be formulated as a rigorous quantitative question. In particular, we believe understanding of the simple repression motif has advanced sufficiently that the design of new versions of the architecture is now possible, based upon predictions about how repressor copy number and DNA binding site strength control expression. In our view, the next step in the progression is to first perform similar rigorous analyses of the fundamental “basis set” of regulatory architectures. Natural follow-ups to this work are explorations of motifs such as simple activation that is regulated by a single activator binding site, and the repressor-activator architecture, mediated by the binding of both a single activator and a single repressor, and beyond. With the individual input-output functions in hand, similar quantitative dissections including the rigorous analysis of their tuning parameters can be undertaken for the “basis set” of full gene-regulatory networks such as switches, feed-forward architectures and oscillators for example, building upon the recent impressive bonanza of efforts from systems biologists and synthetic biologists [74, 75].

References

- ¹F. Jacob and J. Monod, “Genetic regulatory mechanisms in the synthesis of proteins”, *Journal of Molecular Biology* **3**, 318–356 (1961).
- ²R. J. Britten and E. H. Davidson, “Gene regulation for higher cells: a theory”, *Science* **165**, 349–57 (1969).
- ³S. B.-T. de-Leon and E. H. Davidson, “Gene regulation: gene control network in development”, *Annual Review of Biophysics and Biomolecular Structure* **36**, 191 (2007).
- ⁴M. Rydenfelt, H. G. Garcia, R. S. Cox III, and R. Phillips, “The influence of promoter architectures and regulatory motifs on gene expression in *Escherichia coli*”, *PLoS One* **9**, e114347 (2014).
- ⁵N. M. Belliveau, S. L. Barnes, W. T. Ireland, D. L. Jones, M. J. Sweredoski, A. Moradian, S. Hess, J. B. Kinney, and R. Phillips, “Systematic approach for dissecting the molecular mechanisms of transcriptional regulation in bacteria”,

- Proceedings of the National Academy of Sciences of the United States of America **115**, E4796–E4805 (2018).
- ⁶S. Ghatak, Z. A. King, A. Sastry, and B. O. Palsson, “The y-ome defines the 35% of *Escherichia coli* genes that lack experimental evidence of function”, *Nucleic Acids Research* **47**, 2446–2454 (2019).
- ⁷A. Santos-Zavaleta et al., “RegulonDB v 10.5: tackling challenges to unify classic and high throughput knowledge of gene regulation in *E. coli K-12*”, *Nucleic Acids Research* **47**, 212–220 (2019).
- ⁸G. K. Ackers, A. D. Johnson, and M. A. Shea, “Quantitative model for gene regulation by lambda phage repressor”, *Proceedings of The National Academy Of Sciences Of The United States Of America* **79**, 1129–33 (1982).
- ⁹M. A. Shea and G. K. Ackers, “The OR control system of bacteriophage lambda. A physical-chemical model for gene regulation”, *Journal of Molecular Biology* **181**, 211–30 (1985).
- ¹⁰N. E. Buchler, U. Gerland, and T. Hwa, “On schemes of combinatorial transcription logic.”, *Proceedings of the National Academy of Sciences* **100**, 5136–5141 (2003).
- ¹¹J. M. Vilar and S. Leibler, “DNA looping and physical constraints on transcription regulation”, *Journal of Molecular Biology* **331**, 981–9 (2003).
- ¹²J. M. Vilar, C. C. Guet, and S. Leibler, “Modeling network dynamics: the lac operon, a case study”, *Journal of Cell Biology* **161**, 471–6 (2003).
- ¹³L. Bintu, N. E. Buchler, H. G. Garcia, U. Gerland, T. Hwa, J. Kondev, and R. Phillips, “Transcriptional regulation by the numbers: models”, *Current Opinion in Genetics & Development* **15**, 116–24 (2005).
- ¹⁴L. Bintu, N. E. Buchler, H. G. Garcia, U. Gerland, T. Hwa, J. Kondev, T. Kuhlman, and R. Phillips, “Transcriptional regulation by the numbers: applications”, *Current Opinion in Genetics & Development* **15**, 125–135 (2005).
- ¹⁵J. Gertz, E. D. Siggia, and B. A. Cohen, “Analysis of combinatorial cis-regulation in synthetic and genomic promoters”, *Nature* **457**, 215–8 (2009).
- ¹⁶M. S. Sherman and B. A. Cohen, “Thermodynamic state ensemble models of cis-Regulation”, *PLoS computational biology* **8**, e1002407 (2012).
- ¹⁷J. M. Vilar and L. Saiz, “Reliable prediction of complex phenotypes from a modular design in free energy space: an extensive exploration of the lac operon”, *ACS Synth Biol* **2**, 576–86 (2013).
- ¹⁸M. S. Ko, “A stochastic model for gene induction”, *Journal of Theoretical Biology* **153**, 181–94 (1991).
- ¹⁹J. Peccoud and B. Ycart, “Markovian modeling of gene product synthesis”, *Theoretical Population Biology* **48**, 222–234 (1995).

- ²⁰M. T. J. Record, W. Reznikoff, M. Craig, K. McQuade, and P. Schlax, “Escherichia coli RNA polymerase (sigma70) promoters and the kinetics of the steps of transcription initiation”, in *In Escherichia coli and Salmonella cellular and molecular biology*, edited by F. C. Neidhardt et al. (ASM Press, Washington, DC, 1996), pp. 792–821.
- ²¹T. B. Kepler and T. C. Elston, “Stochasticity in transcriptional regulation: origins, consequences, and mathematical representations.”, *Biophys J.* **81**, 3116–36. (2001).
- ²²A. Sanchez and J. Kondev, “Transcriptional control of noise in gene expression”, *Proceedings of The National Academy Of Sciences Of The United States Of America* **105**, 5081–6 (2008).
- ²³V. Shahrezaei and P. S. Swain, “Analytical distributions for stochastic gene expression”, *Proceedings of the National Academy of Sciences* **105**, 17256–17261 (2008).
- ²⁴A. Sanchez, H. G. Garcia, D. Jones, R. Phillips, and J. Kondev, “Effect of promoter architecture on the cell-to-cell variability in gene expression”, *PLoS computational biology* **7**, e1001100 (2011).
- ²⁵D. Michel, “How transcription factors can adjust the gene expression floodgates”, *Progress in Biophysics and Molecular Biology* **102**, 16–37 (2010).
- ²⁶H. G. Garcia, J. Kondev, N. Orme, J. A. Theriot, and R. Phillips, “Thermodynamics of Biological Processes”, in *Methods in Enzymology*, Vol. 492 (Elsevier, 2011), pp. 27–59.
- ²⁷R. Phillips, N. M. Belliveau, G. Chure, H. G. Garcia, M. Razo-Mejia, and C. Scholes, “Figure 1 Theory Meets Figure 2 Experiments in the Study of Gene Expression”, *Annual Review of Biophysics* **48**, 121–163 (2019).
- ²⁸W. T. Ireland, S. M. Beeler, E. Flores-Bautista, N. M. Belliveau, M. J. Sweredoski, A. Moradian, J. B. Kinney, and R. Phillips, *Deciphering the regulatory genome of Escherichia coli , one hundred promoters at a time*, Preprint (Biophysics, Jan. 2020).
- ²⁹S. Oehler, M. Amouyal, P. Kolkhof, B. von Wilcken-Bergmann, and B. Müller-Hill, “Quality and position of the three *lac* operators of *E. coli* define efficiency of repression”, *The EMBO journal* **13**, 3348–3355 (1994).
- ³⁰S. Oehler, E. R. Eismann, H. Krämer, and B. Müller-Hill, “The three operators of the *lac* operon cooperate in repression”, *The EMBO journal* **9**, 973–979 (1990).
- ³¹S. A. Frank, “Input-output relations in biological systems: measurement, information and the Hill equation”, *Biology Direct* **8**, 31 (2013).
- ³²M. Razo-Mejia, S. L. Barnes, N. M. Belliveau, G. Chure, T. Einav, M. Lewis, and R. Phillips, “Tuning Transcriptional Regulation through Signaling: A Predictive Theory of Allosteric Induction”, *Cell Systems* **6**, 456–469.e10 (2018).

- ³³H. G. Garcia and R. Phillips, “Quantitative dissection of the simple repression input-output function”, *Proceedings of the National Academy of Sciences* **108**, 12173–12178 (2011).
- ³⁴R. C. Brewster, F. M. Weinert, H. G. Garcia, D. Song, M. Rydenfelt, and R. Phillips, “The Transcription Factor Titration Effect Dictates Level of Gene Expression”, *Cell* **156**, 1312–1323 (2014).
- ³⁵G. Chure, M. Razo-Mejia, N. M. Belliveau, T. Einav, Z. A. Kaczmarek, S. L. Barnes, M. Lewis, and R. Phillips, “Predictive shifts in free energy couple mutations to their phenotypic consequences”, *Proceedings of the National Academy of Sciences* **116**, 18275–18284 (2019).
- ³⁶D. L. Jones, R. C. Brewster, and R. Phillips, “Promoter architecture dictates cell-to-cell variability in gene expression”, *Science* **346**, 1533–1536 (2014).
- ³⁷R. Phillips, J. Kondev, J. Theriot, and H. G. Garcia, *Physical biology of the cell, 2nd Edition* (Garland Science, New York, 2013).
- ³⁸N. Mitarai, S. Semsey, and K. Sneppen, “Dynamic competition between transcription initiation and repression: Role of nonequilibrium steps in cell-to-cell heterogeneity”, *Physical Review E* **92**, 022710 (2015).
- ³⁹P. L. deHaseth, M. L. Zupancic, and M. T. Record, “RNA Polymerase-Promoter Interactions: the Comings and Goings of RNA Polymerase”, *Journal of Bacteriology* **180**, 3019–3025 (1998).
- ⁴⁰E. L. King and C. Altman, “A Schematic Method of Deriving the Rate Laws for Enzyme-Catalyzed Reactions”, *The Journal of Physical Chemistry* **60**, 1375–1378 (1956).
- ⁴¹T. L. Hill, “Studies in irreversible thermodynamics IV. diagrammatic representation of steady state fluxes for unimolecular systems”, *Journal of Theoretical Biology* **10**, 442–459 (1966).
- ⁴²M. Razo-Mejia, S. Marzen, G. Chure, R. Taubman, M. Morrison, and R. Phillips, “First-principles prediction of the information processing capacity of a simple genetic circuit”, arXiv:2005.03214 [q-bio] (2020).
- ⁴³S. A. Frank, “Generative models versus underlying symmetries to explain biological pattern”, *Journal of Evolutionary Biology* **27**, 1172–1178 (2014).
- ⁴⁴J. Gunawardena, “A Linear Framework for Time-Scale Separation in Nonlinear Biochemical Systems”, *PLoS ONE* **7**, edited by K. Selvarajoo, e36321 (2012).
- ⁴⁵T. Ahsendorf, F. Wong, R. Eils, and J. Gunawardena, “A framework for modelling gene regulation which accommodates non-equilibrium mechanisms”, *BMC Biology* **12**, 102 (2014).
- ⁴⁶T. L. Hill, *Free Energy Transduction and Biochemical Cycle Kinetics* (Springer New York, New York, NY, 1989).

- ⁴⁷S. Chong, C. Chen, H. Ge, and X. S. Xie, “Mechanism of Transcriptional Bursting in Bacteria”, *Cell* **158**, 314–326 (2014).
- ⁴⁸S. A. Sevier, D. A. Kessler, and H. Levine, “Mechanical bounds to transcriptional noise”, *Proceedings of the National Academy of Sciences* **113**, 13983–13988 (2016).
- ⁴⁹I. I. Cisse et al., “Real-Time Dynamics of RNA Polymerase II Clustering in Live Human Cells”, *Science* **341**, 664–667 (2013).
- ⁵⁰W.-K. Cho et al., “RNA Polymerase II cluster dynamics predict mRNA output in living cells”, *eLife* **5**, e13617 (2016).
- ⁵¹A.-M. Ladouceur et al., “Clusters of bacterial RNA polymerase are biomolecular condensates that assemble through liquid-liquid phase separation”, *10.1101/2020.03.16.994491* (2020).
- ⁵²A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*, Third edition, Chapman & Hall/CRC Texts in Statistical Science (CRC Press, Boca Raton, 2013).
- ⁵³J. B. Kinney, A. Murugan, C. G. Callan, and E. C. Cox, “Using deep sequencing to characterize the biophysical mechanism of a transcriptional regulatory sequence”, *Proceedings of the National Academy of Sciences* **107**, 9158–9163 (2010).
- ⁵⁴R. C. Brewster, D. L. Jones, and R. Phillips, “Tuning Promoter Strength through RNA Polymerase Binding Site Design in *Escherichia coli*”, *PLoS Computational Biology* **8**, edited by E. van Nimwegen, e1002811 (2012).
- ⁵⁵J. Landman, R. N. Georgiev, M. Rydenfelt, and W. K. Kegel, “*IN VIVO* and *in vitro* consistency of thermodynamic models for transcription regulation”, *Physical Review Research* **1**, 033094 (2019).
- ⁵⁶P. Hammar, M. Walldén, D. Fange, F. Persson, Ö. Baltekin, G. Ullman, P. Leroy, and J. Elf, “Direct measurement of transcription factor dissociation excludes a simple operator occupancy model for gene regulation”, *Nature Genetics* **46**, 405–408 (2014).
- ⁵⁷A. Sanchez, S. Choubey, and J. Kondev, “Stochastic models of transcription: From single molecules to single cells”, *Methods* **62**, 13–25 (2013).
- ⁵⁸R. Phillips, “Napoleon Is in Equilibrium”, *Annual Review of Condensed Matter Physics* **6**, 85–111 (2015).
- ⁵⁹L.-h. So, A. Ghosh, C. Zong, L. A. Sepúlveda, R. Segev, and I. Golding, “General properties of transcriptional time series in *Escherichia coli*”, *Nature Genetics* **43**, 554–560 (2011).
- ⁶⁰I. Golding, J. Paulsson, S. M. Zawilski, and E. C. Cox, “Real-Time Kinetics of Gene Activity in Individual Bacteria”, *Cell* **123**, 1025–1036 (2005).
- ⁶¹J. K. Blitzstein and J. Hwang, *Introduction to probability*, Texts in Statistical Science (CRC Press/Taylor & Francis Group, Boca Raton, 2015).

- ⁶²M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, “Perspective: Sloppiness and emergent theories in physics, biology, and beyond”, *The Journal of Chemical Physics* **143**, 010901 (2015).
- ⁶³G. Chure, Z. A. Kaczmarek, and R. Phillips, “Physiological Adaptability and Parametric Versatility in a Simple Genetic Circuit”, *bioRxiv*, 10.1101/2019.12.19.878462 (2019).
- ⁶⁴H. Chen, K. Shiroguchi, H. Ge, and X. S. Xie, “Genome-wide study of mRNA degradation and transcript elongation in *Escherichia coli*”, *Molecular Systems Biology* **11**, 781 (2015).
- ⁶⁵E. T. Jaynes and G. L. Bretthorst, *Probability theory: the logic of science* (Cambridge University Press, Cambridge, UK ; New York, NY, 2003).
- ⁶⁶A. Gelman and C. R. Shalizi, “Philosophy and the practice of Bayesian statistics: *Philosophy and the practice of Bayesian statistics*”, *British Journal of Mathematical and Statistical Psychology* **66**, 8–38 (2013).
- ⁶⁷J. A. Englaender, J. A. Jones, B. F. Cress, T. E. Kuhlman, R. J. Linhardt, and M. A. G. Koffas, “Effect of Genomic Integration Location on Heterologous Protein Expression and Metabolic Engineering in *E. coli*”, *ACS Synthetic Biology* **6**, 710–720 (2017).
- ⁶⁸C. Engl, G. Jovanovic, R. D. Brackston, I. Kotta-Loizou, and M. Buck, “The route to transcription initiation determines the mode of transcriptional bursting in *E. coli*”, *Nature Communications* **11**, 2422 (2020).
- ⁶⁹T. M. Dunn, S. Hahn, S. Ogden, and R. F. Schleif, “An operator at -280 base pairs that is required for repression of araBAD operon promoter: addition of DNA helical turns between the operator and promoter cyclically hinders repression”, *Proceedings of the National Academy of Sciences of the United States of America* **81**, 5017–5020 (1984).
- ⁷⁰M. J. Weickert and S. Adhya, “The galactose regulon of *Escherichia coli*”, *Molecular Microbiology* **10**, 245–251 (1993).
- ⁷¹R. Schleif, “Regulation of the L-arabinose operon of *Escherichia coli*”, *Trends in Genetics* **16**, 559–565 (2000).
- ⁷²S. Semsey, M. Geanakopoulos, D. E. A. Lewis, and S. Adhya, “Operator-bound GalR dimers close DNA loops by direct interaction: tetramerization and inducer binding”, *The EMBO journal* **21**, 4349–4356 (2002).
- ⁷³L. Swint-Kruse and K. S. Matthews, “Allostery in the LacI/GalR family: variations on a theme”, *Current Opinion in Microbiology* **12**, 129–137 (2009).
- ⁷⁴R. Milo, “Network motifs: simple building blocks of complex networks”, *Science* **298**, 824–827 (2002).

- ⁷⁵U. Alon, *An introduction to systems biology: design principles of biological circuits*, Chapman & Hall/CRC Mathematical and Computational Biology Series 10 (Chapman & Hall/CRC, Boca Raton, FL, 2007).

Chapter 3

THEORETICAL INVESTIGATION OF A GENETIC SWITCH FOR METABOLIC ADAPTATION

¹K. S. Laxhuber, M. J. Morrison, G. Chure, N. M. Belliveau, C. Strandkvist, K. L. Naughton, and R. Phillips, “Theoretical investigation of a genetic switch for metabolic adaptation”, PLOS ONE **15**, edited by L. M. Espinoza-Fonseca, e0226453 (2020).

3.1 Abstract

Membrane transporters carry key metabolites across the cell membrane and, from a resource standpoint, are hypothesized to be produced when necessary. The expression of membrane transporters in metabolic pathways is often upregulated by the transporter substrate. In *E. coli*, such systems include for example the *lacY*, *araFGH*, and *xylFGH* genes, which encode for lactose, arabinose, and xylose transporters, respectively. As a case study of a minimal system, we build a generalizable physical model of the *xapABR* genetic circuit, which features a regulatory feedback loop via membrane transport (positive feedback) and enzymatic degradation (negative feedback) of an inducer. Dynamical systems analysis and stochastic simulations show that the membrane transport makes the model system bistable in certain parameter regimes. Thus, it serves as a genetic “on-off” switch, enabling the cell to only produce a set of metabolic enzymes when the corresponding metabolite is present in large amounts. We find that the negative feedback from the degradation enzyme does not significantly disturb the positive feedback from the membrane transporter. We investigate hysteresis in the switching and discuss the role of cooperativity and multiple binding sites in the model circuit. Fundamentally, this work explores how a stable genetic switch for a set of enzymes is obtained from transcriptional auto-activation of a membrane transporter through its substrate.

3.2 Introduction

Genetic regulatory circuits are fundamental building blocks of functioning cells and organisms. One abundant class of these circuits are genetic switches. Although their construction and function may differ, their common feature is bistability: their output gene expression will flow to and remain at one of two steady-state levels.




Figure 3.1: **A schematic of different genetic switches.** (A) and (B) show the two most well-known genetic switches: (A) two mutual repressors and (B) a self-activating gene. In (C), a very much simplified version of the circuit that we investigate in this paper can be seen, where the similarity to the switch in (B) is clear. A complete version of the model circuit can be found in Fig 3.3.

The distribution of gene expression in a cell culture can then be bimodal. This is not to be confused with mere stochastic bimodality, where the system is not stable, and the gene expression in each cell can fluctuate between the two levels.

One classic example of a genetic switch is a system where two repressor proteins each regulate the transcription of the other [1, 2] (illustrated schematically in Fig 3.1). Here, one stable state is high expression of the first protein and low expression of the second, and the second stable state is the opposite. This switch enables the system to have a memory: if something induces expression of either one of the proteins, the system will remain in this state until a significant perturbation occurs. Another well-known and even simpler example is an auto-activating circuit in which a protein activates its own transcription [3]. This gives the system an “on-off” switch.

Through physical and mathematical modeling, we investigate a more complex switch system where the bistability is due, as we will show, to a membrane transport protein. Such a switch is common for metabolic processes in biology, for reasons discussed below. Existing models in the literature tend towards one of two extremes: either highly detailed descriptions of specific, complicated networks (e.g., [4]), or Hill function descriptions that coarse-grain all complexity into a few parameters with inscrutable microscopic physical meaning. We aim for a middle ground in this work. We seek an intuitive understanding through a simple model of a minimal system, with only the essential components and interactions for the questions we pose. Yet we still model these components explicitly and discuss the necessary model complexity for a physically correct model.

The key feature of the type of system we investigate is the indirect activation of the

transporter gene by the transporter substrate, leading to positive feedback similar to the aforementioned “on-off” switch. An example for such an architecture is the *lac* operon, where lactose indirectly activates the expression of lactose permease. Other examples in *E. coli* include the *araFGH* and *xylFGH* operons, which contain genes for arabinose and xylose transporters, respectively. For *lac* and *araFGH*, bistability has indeed been observed and attributed to such a positive feedback loop, for example in the well-known study by Novick and Weiner (1957), among other works [5–12]. A eukaryotic example is the glucose transporter GLUT-2 in liver and β -cells [13, 14], though this system is much more complex than the following analysis.

It is quite conceivable that this auto-activation process is common to many substances that a cell would want to consume. Such a switch enables the cell to sense and respond to its environment: if the substrate enters the cell, it activates the production of membrane transporters. The cell then starts accumulating the substrate, thereby “testing” the substrate’s presence in the extracellular environment. If there is enough, the expression stabilizes at an “on” state and the cell has, in a short-term sense, adapted. When there is not enough substrate, the operon, which often encodes for a whole set of enzymes for this one metabolite, switches “off” again. Such a mechanism could plausibly be involved in various cases of short-term adaptation such as, but not limited to, the *lacZYA*, *araFGH*, and *xylFGH* examples mentioned above.

A key element of this mechanism is the presence of a transcription factor which binds to the transporter substrate and which is often expressed at a low level (often at copy numbers of order ~ 10 , [15, 16]). This is resource efficient for the cell, as this low copy number transcription factor acts as an “always on” sensor to detect the substrate, allowing high copy numbers of the membrane transporter and its attendant operon to be expressed only when their substrate is actually present. The transcription factors LacI, AraC, and XylR all appear to fill this role [5–12, 15, 17].

For our modeling, we focus on the *xapABR* genetic circuit from *E. coli* as a case study. It is similar to *lac*, but less complex. Instead of lactose, its purpose is to make use of the nucleoside xanthosine as an energy source [18, 19]. The circuit is made up of two operons: one that encodes for XapR and another that encodes for XapA and XapB. XapR is a transcription factor that is induced by xanthosine and activates the *xapAB* promoter, in close analogy to AraC, XylR, and also LacI.¹ The

¹One might object that LacI represses its target operon, while XapR, AraC, and XylR activate

xapAB promoter has been suggested to have two binding sites for XapR[20], but the promoter architecture and function is not yet fully understood. The transcription of *xapR* seems to be constitutive and not auto-regulated [20]. Structural homology to other transcription factors suggests that XapR appears in dimers where one dimer can bind two xanthosine molecules [21]. The protein XapA is a purine nucleoside phosphorylase that degrades xanthosine into components (ribose and xanthine) that can be fed into metabolic pathways [18, 19]. XapB on the other hand is a membrane transporter of xanthosine [20, 22].

Experimentally, we found that the expression level of *xapAB* among cells is bimodal and that the system seems to be bistable (see next section). We aim to understand which of the circuit’s features are necessary for bistability and investigate its behavior in different parameter regimes. After presenting some experimental background on the *xap* genes, we discuss the details of our model. Lastly, we estimate the free parameters and then present the observations we made through phase diagrams, followed by the results from stochastic simulations.

3.3 Experimental motivation

Our work was motivated by the experimental observation of bimodality in the *xap* circuit, which is shown in Fig 3.2. We focus on the essential findings here and refer the reader to B.1 for more experimental details. Briefly, we placed a fluorescent reporter under the control of the wild type (wt) *xapAB* promoter. This construct was placed in three different backgrounds: $\Delta xapABR$, $\Delta xapAB$, and wt, and expression as a function of extracellular xanthosine concentration was measured using flow cytometry, as has been described previously [23]. The left panel of Figure 3.2A shows that for increasing xanthosine concentration, expression in the wt-background increases in a switch-like way: it is nearly zero when there is little xanthosine (“uninduced state”), but increases drastically when there is more (“induced state”). In between, the aforementioned bimodal distribution is obtained, where some bacteria are in the uninduced expression state and others are induced.

When all genes of the *xapABR* circuit are removed, the xanthosine response of our reporter construct disappears (see Figure 3.2A, middle panel). This is the result of removing the transcription factor XapR that is induced by xanthosine. When only the genes *xapAB* are removed but *xapR* is kept, a response to the xanthosine

their target operons. However, the analogy we wish to draw is that the qualitative logic of their inducers are all identical, i.e., the presence of their respective inducer causes their target operon to be transcribed.

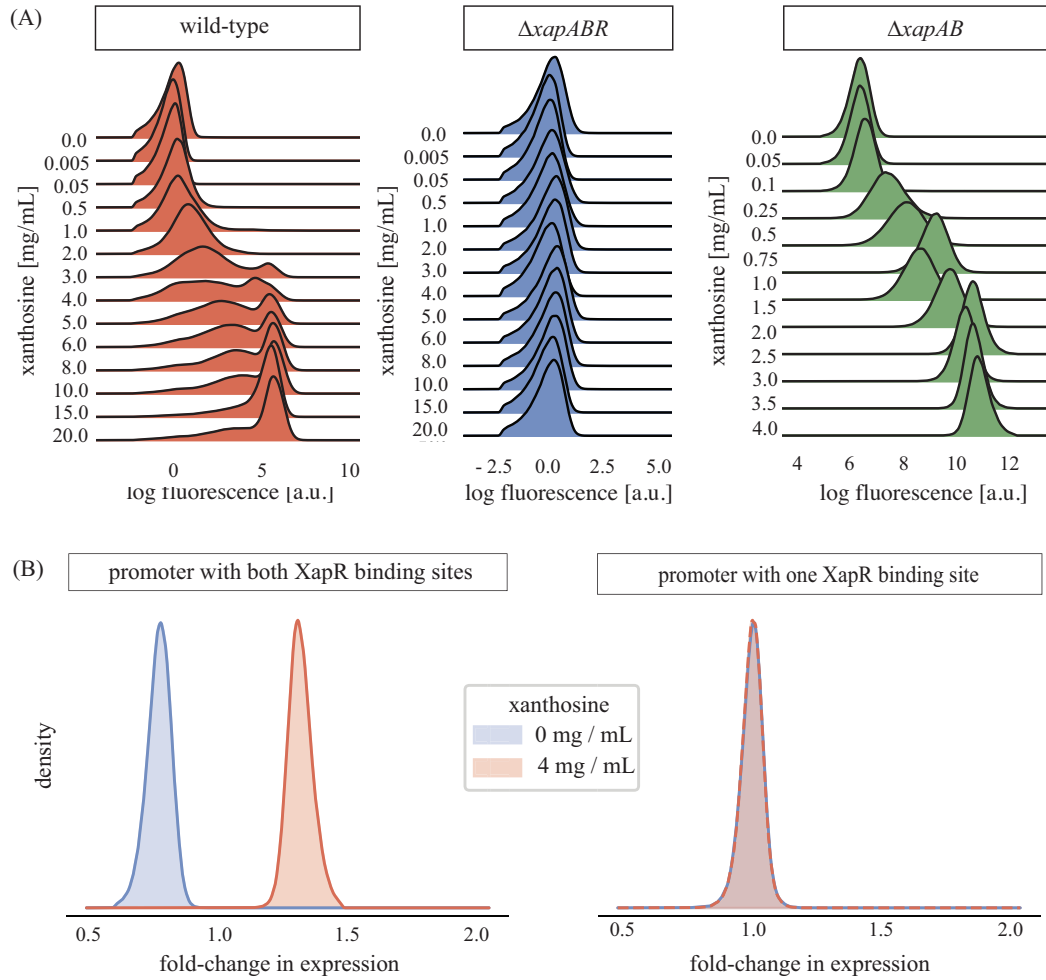


Figure 3.2: Experimental data on the *xap* circuit. (A) The expression of the *xapAB* promoter was measured for different extracellular concentrations of xanthosine (vertical axis). The left panel shows the wild type circuit while the middle and right panels show the effect of deleting the genes *xapABR* and *xapAB*, respectively. The wild type circuit behaves like a switch. Note that the fluorescence scale of the middle panel is not comparable with the other two, and also that the chosen xanthosine concentrations are different. (B) shows the fold-change in gene expression upon addition of xanthosine for the wt promoter (left panel) and for the promoter with only the XapR binding site adjacent to the polymerase binding site (right panel). Note that the fold-change used here differs from fold-change in, e.g., [23–25], in that no subtraction of autofluorescence was performed, which is adequate for the qualitative comparison of these two promoters.

concentration is regained but it is no longer switch-like (see Figure 3.2A, right panel). Instead, the distribution remains unimodal and simply shifts to higher expression as the xanthosine concentration is increased. As we will see later on, this is because the circuit now lacks the positive feedback loop due to the xanthosine membrane transporter XapB.

As mentioned in the introduction, there are two binding sites for XapR on the *xapAB* promoter: one partially overlapping the polymerase -35 site and one further upstream. Working in the $\Delta xapAB$ background, we measured the expression level of our reporter when driven by two constructs: the wild-type *xapAB* promoter (Figure 3.2B, left panel), and the *xapAB* promoter with the upstream binding site removed (Figure 3.2B, right panel). Clearly, removing one XapR binding site dramatically reduced the responsiveness of the promoter to xanthosine (though not apparent from Figure 3.2B, this response remains detectable in other measurements covering a smaller dynamic range).

Already, we believe these data, combined with the prior knowledge of the field, clearly suggest a minimal model of the *xapABR* system. The construction and exploration of this model occupies the remainder of this work. We leave it for future work to quantitatively dissect and test the model, in the manner of [23–25], for example.

3.4 Model

Step by step modeling of the system

In this section, we present our model of the *xapABR* genetic circuit. Fig 3.3 shows an overview of this model. The qualitative picture of the circuit switching its state is as follows:

- In the initial absence of XapB, small amounts of xanthosine permeate into and out of the cell (discussed in more detail below).
- The presence of xanthosine shifts XapR's equilibrium from inactive to active.
- Active XapR binds to the *xapAB* promoter, increasing transcription.
- From this mRNA transcript, translation produces the two proteins XapA and XapB.
- XapB actively transports much larger amounts of xanthosine into the cell and XapA degrades xanthosine.




Figure 3.3: **Model of the *xapAB* circuit.** The XapR dimers are induced by xanthosine and the induced XapR binds cooperatively as an activator to the *xapAB* promoter. For these two steps, quasi-equilibrium is assumed. If both XapR binding sites are occupied and the polymerase is bound, the gene is transcribed at rate r_m . The mRNA decays at rate γ_m , and both proteins are translated at rate r_p and decay at rate γ_p . XapA degrades xanthosine with Michaelis-Menten parameters k_a and K_a . Similarly, XapB is treated as a Michaelis-Menten enzyme which imports ($k_{b,i}$, $K_{b,i}$) and exports ($k_{b,e}$, $K_{b,e}$) xanthosine. Furthermore, xanthosine enters and leaves the cell through non-specific transport, proportional to rates k_{nup} and ξk_{nup} , respectively.

- Production of more XapA and XapB is balanced by degradation or dilution through cell divisions.

Because xanthosine activates the transcription factor XapR, we have positive and negative feedback loops due to XapB and XapA, respectively. The remainder of this subsection discusses each of the above steps in detail, leading us to a set of two coupled ODEs. More in-depth explanations can be found in B.1.

Induction of XapR. We treat dimers as the only form of XapR that appears in the cell. Each dimer can bind two xanthosine molecules [21]. The Monod-Wyman-Changeux (MWC) model is used to describe the fraction of XapR dimers in the

active state, which has the form

$$[\text{XapR}]_A = [\text{XapR}]_{\text{tot}} \frac{\left(1 + \frac{[x]}{K_{xA}}\right)^2}{\left(1 + \frac{[x]}{K_{xA}}\right)^2 + e^{\beta \Delta E_x} \left(1 + \frac{[x]}{K_{xA}} \frac{K_{xA}}{K_{xI}}\right)^2}. \quad (3.1)$$

A detailed discussion of the MWC model can be found in [26]. In Eq 3.1, $[x]$ is the xanthosine concentration, and $[\text{XapR}]_A$ and $[\text{XapR}]_{\text{tot}}$ denote the concentration of active and total XapR dimers, respectively. Furthermore, K_{xI} and K_{xA} are the dissociation constants of xanthosine to the inactive and the active XapR dimer, respectively, and ΔE_x stands for the energy difference between the inactive and the active states of the protein. We expect $\Delta E_x > 0$ and $K_{xA} < K_{xI}$ for inducible activation. This corresponds to XapR being mainly inactive in the absence of xanthosine and becoming mostly active at high concentrations of xanthosine.

Transcription. Transcription and translation of the *xapAB* operon, regulated by the induced XapR, produce the two proteins XapA and XapB. We start with transcription and assume that the binding of XapR and polymerase to the promoter is at quasi-equilibrium. The polymerase binding is modeled as independent of that of XapR, and all influence of the activator is pushed into the transcription rate. Furthermore, the binding energy of XapR to each of its two sites is assumed to be the same. A discussion of these simplifications can be found in B.1.

In Fig 3.4, all possible states of the promoter in our model and their corresponding thermodynamic weights are shown. $[P]$ denotes the polymerase concentration, and ΔE_{coop} stands for the interaction energy of the two XapR dimers. If cooperativity in transcription factor binding is neglected, this is set to zero. Furthermore, K_{XapR} and K_P denote the dissociation constant of XapR and polymerase to the promoter, respectively. In statistical mechanics language these dissociation constants are equivalent to $\frac{N_{\text{NS}}}{V} e^{\beta \Delta E_{\text{XapR}}}$ and $\frac{N_{\text{NS}}}{V} e^{\beta \Delta E_P}$, respectively, with N_{NS} being the number of non-specific binding sites on the DNA, V the volume of the cell, and ΔE_{XapR} and ΔE_P , respectively, the interaction energies of XapR or polymerase with the promoter.

We consider only the state where both XapR binding sites are occupied as active and all other states as inactive, meaning they have transcription rate equal to zero. Our experiments show that the expression becomes very weak when one of the XapR binding sites is removed from the promoter (see Figure 3.2B and B.1), suggesting that this simplification is reasonable. Furthermore, we find that in the bistable

	State	Weight
①		1
②		$\frac{[\text{XapR}]_A}{K_{\text{XapR}}}$
③		$\frac{[\text{XapR}]_A}{K_{\text{XapR}}}$
④		$\left(\frac{[\text{XapR}]_A}{K_{\text{XapR}}}\right)^2 e^{-\beta \Delta E_{\text{coop}}}$
⑤		$\frac{[\text{P}]}{K_P}$
⑥		$\frac{[\text{P}]}{K_P} \frac{[\text{XapR}]_A}{K_{\text{XapR}}}$
⑦		$\frac{[\text{P}]}{K_P} \frac{[\text{XapR}]_A}{K_{\text{XapR}}}$
⑧		$\frac{[\text{P}]}{K_P} \left(\frac{[\text{XapR}]_A}{K_{\text{XapR}}}\right)^2 e^{-\beta \Delta E_{\text{coop}}}$

Figure 3.4: **The promoter states.** We consider only the completely occupied state as active and all other states (faded out in the figure) as completely inactive. The parameters are the interaction energy of the two XapR dimers ΔE_{coop} and the dissociation constants K_{XapR} and K_P of XapR and polymerase to the promoter, respectively. The concentrations of polymerase and active XapR are denoted by $[\text{P}]$ and $[\text{XapR}]$.

parameter range, considering the single occupancy states as active instead has almost no influence on the results (see also B.1).

With $[\text{m}]$ being the mRNA concentration, r_m the transcription rate, γ_m the mRNA decay rate, and p_{active} the probability of the promoter being in the active state, we obtain

$$\frac{d[\text{m}]}{dt} = r_m p_{\text{active}} - \gamma_m [\text{m}] \quad (3.2)$$

$$p_{\text{active}} = \frac{w_8}{\sum_{i=1}^8 w_i} = \frac{[\text{P}]}{K_P + [\text{P}]} \frac{\left(\frac{[\text{XapR}]_A}{K_{\text{XapR}}}\right)^2 e^{-\beta \Delta E_{\text{coop}}}}{1 + 2 \frac{[\text{XapR}]_A}{K_{\text{XapR}}} + \left(\frac{[\text{XapR}]_A}{K_{\text{XapR}}}\right)^2 e^{-\beta \Delta E_{\text{coop}}}} \quad (3.3)$$

Here, w_i stands for the thermodynamic weight of the i th state in the order in which they are listed in Fig 3.4. As written above, the partition function factorizes into a polymerase and a XapR term because of our assumption of independent binding, which is further discussed in B.1. Note that because r_m implicitly contains the gene copy number per cell, it has units of $\text{M}^{-1} \text{s}^{-1}$ and not just s^{-1} . This rate equation gives the mean mRNA concentration $\langle [\text{mRNA}] \rangle = \frac{r_m}{\gamma_m} p_{\text{active}}$, which we will need in the next paragraph. The mean can also be found from the full chemical master equation, which is provided in B.1.

Translation. The next step in our modeling progression is translation. As a simplification, we write $[p] = [\text{XapA}] = [\text{XapB}]$ for the general protein concentration. This assumes that the rates of transcription, mRNA decay, translation, and protein decay are the same for both proteins, which, as discussed in B.1, does not have a significant influence on the results. We write the following rate equation for the protein concentration, where r_p denotes the translation rate, γ_p the protein decay rate, and $\langle[m]\rangle$ the mean mRNA concentration:

$$\frac{d[p]}{dt} = r_p \langle[m]\rangle - \gamma_p [p]. \quad (3.4)$$

Xanthosine dynamics. Having described how xanthosine activates the synthesis of XapA and XapB through XapR, we now close the feedback loop by setting up a xanthosine rate equation.

There are two significant mechanisms for transport of xanthosine across the cell membrane. In the induced system, the main transporter is XapB, whereas in the uninduced system, there is almost no XapB. Instead, xanthosine can enter the cell through the two nucleoside transporters NupC and NupG, which have a very low affinity for xanthosine [22]. All these transporters, XapB, NupC, and NupG, are powered by the proton gradient across the membrane [22], which is why we assume their kinetic scheme to be similar to that of the *lac* permease (as it is described in [27]). There can be import and export of xanthosine, and which one dominates depends on the proton and xanthosine concentrations on the two sides of the membrane. In both cases, a proton and a substrate need to bind to the transporter on one side of the membrane and detach from it on the other side before the empty transporter moves back to the other side again. We refer the reader to B.1 for a detailed description of the transport and its modeling, and just state the result here. We model influx and efflux separately. For XapB, we use Michaelis-Menten kinetics with parameters $k_{b,i}$, $K_{b,i}$ for influx and $k_{b,e}$, $K_{b,e}$ for efflux. For the Nup transporters, we also use Michaelis-Menten kinetics but, because of the transporter's low affinity for xanthosine, we can linearize the Michaelis-Menten term as $k_{\text{cat}}[x]/(K_M + [x]) \approx \tilde{k}[x]$ (i.e., $K_M \gg [x]$ across the physiologically relevant range for $[x]$). For the rate parameters \tilde{k} , we write k_{nup} for influx and ξk_{nup} for efflux.

After transport into the cell, XapA degrades xanthosine. We model this using standard Michaelis-Menten kinetics, with parameters k_a , K_a (corresponding to turnover rate and Michaelis constant, respectively). Transport and degradation then leads to

the xanthosine rate equation

$$\frac{d[x]}{dt} = \underbrace{\left(k_{b,i} \frac{c}{K_{b,i} + c} - k_{b,e} \frac{[x]}{K_{b,e} + [x]} \right)}_{\text{XapB}} - \underbrace{k_a \frac{[x]}{K_a + [x]}}_{\text{XapA}} [p] + \underbrace{k_{\text{nup}} (c - \xi[x])}_{\text{NupC \& NupG}}. \quad (3.5)$$

Recall that $[x]$ is the intracellular xanthosine concentration, while c denotes the extracellular concentration. Because $k_{b,i} > k_{b,e}$ and $K_{b,i} < K_{b,e}$, influx dominates at low intracellular xanthosine concentrations. At much higher intra- than extracellular xanthosine concentrations, the efflux term takes over. More details on the aforementioned steps and a discussion of passive diffusion can be found in B.1.

Nondimensionalization

We have now formulated the behavior of the system in terms of the rate equations for mRNA, protein, and xanthosine. These equations can be nondimensionalized, which reduces the dimension of parameter space. We measure time in units of γ_p^{-1} and concentrations in units of K_a (except XapR, where the equations make it more natural to use K_{XapR}). In Table 3.1, all the nondimensional parameters and their definitions are listed. Furthermore, we define $[m]_a := \frac{[m]}{K_a}$, $[p]_a := \frac{[p]}{K_a}$, $[x]_a := \frac{[x]}{K_a}$, $[c]_a := \frac{[c]}{K_a}$, and $\tau := \gamma_p t$. Using these definitions, the following equations are obtained

$$\frac{d[m]_a}{d\tau} = \rho_m \frac{[\text{XapR}]_{R,A}^2 e^{-\Delta\epsilon_{\text{coop}}}}{1 + 2[\text{XapR}]_{R,A} + [\text{XapR}]_{R,A}^2 e^{-\Delta\epsilon_{\text{coop}}}} - \gamma_{\text{mp}} [m]_a \quad (3.6)$$

$$\frac{d[p]_a}{d\tau} = \rho_p [m]_a - [p]_a \quad (3.7)$$

$$\begin{aligned} \frac{d[x]_a}{d\tau} = & \left(k_{\beta,i} \frac{[c]_a}{K_{\beta,i} + [c]_a} - k_{\beta,e} \frac{[x]_a}{K_{\beta,e} + [x]_a} - k_\alpha \frac{[x]_a}{1 + [x]_a} \right) [p]_a \\ & + k_\eta ([c]_a - \xi[x]_a) \end{aligned} \quad (3.8)$$

$$\text{with } [\text{XapR}]_{R,A} = [\text{XapR}]_R \frac{\left(1 + \frac{[x]_a}{K_{\text{XA}}}\right)^2}{\left(1 + \frac{[x]_a}{K_{\text{XA}}}\right)^2 + e^{\Delta\epsilon_x} \left(1 + \frac{[x]_a}{K_{\text{XA}}} \frac{1}{K_{\text{IA}}}\right)^2}$$

Very little is known about the *xap* system, and thus, there are almost no measured values for the free parameters. Nevertheless, we were able to estimate a reasonable range by using values from similar, well studied systems and by exploiting physical constraints or relations between parameters. The results of these estimates are shown in Table 3.1. They are based on a choice of $\gamma_p = 5 \cdot 10^{-4} \text{ s}^{-1}$ and $K_a = 5 \cdot 10^{-5} \text{ M}$. A detailed derivation can be found in B.1.

Table 3.1: **Nondimensional parameters and their estimated values.**

Param.	Definition	Estimated range	Value used
ρ_m	$:= \frac{r_m}{\gamma_p K_a} \frac{[P]}{K_p + [P]}$	$\approx 10^{-3 \pm 2}$	10^{-3}
γ_{mp}	$:= \frac{\gamma_m}{\gamma_p}$	$\approx 10^{1 \pm 0.5}$	10^1
ρ_p	$:= \frac{r_p}{\gamma_p}$	$\approx 10^{2 \pm 0.5}$	10^2
$[XapR]_R$	$:= \frac{[XapR]_{tot}}{K_{XapR}}$	$\approx 10^{0 \pm 2}$	10^0
$[c]_a$	$:= \frac{c}{K_a}$	$(\in [0, 10^3])$	13
$k_{\beta,i}$	$:= \frac{k_{b,i}}{\gamma_p}$	$\approx 10^{4 \pm 1}$	$5 \cdot 10^4$
$k_{\beta,e}$	$:= \frac{k_{b,e}}{\gamma_p}$	$\approx 10^{3 \pm 2}$	10^3
k_α	$:= \frac{k_a}{\gamma_p}$	$\approx 10^{2 \pm 0.8}$	10^2
k_η	$:= \frac{k_{nup}}{\gamma_p}$	$\approx 10^{0 \pm 3}$	$5 \cdot 10^{-1}$
ξ	$= \xi$	$\approx 0.8 \pm 0.1$	0.8
$K_{\beta,i}$	$:= \frac{K_{b,i}}{K_a}$	$\approx 10^{1 \pm 2}$	10^1
$K_{\beta,e}$	$:= \frac{K_{b,e}}{K_a}$	$\approx 10^{2 \pm 2}$	10^2
$K_{\chi A}$	$:= \frac{K_{\chi A}}{K_a}$	$\approx 10^{2 \pm 1} \cdot 10^{\Delta\epsilon_x - 5}$	10^2
K_{IA}	$:= \frac{K_{xl}}{K_{\chi A}}$	$\approx 10^{2 \pm 1}$	10^2
$\Delta\epsilon_x$	$:= \beta \Delta E_x$	$\approx 2 \text{ to } 2 (\ln(K_{IA}) - 1) < 12$	5
$\Delta\epsilon_{coop}$	$:= \beta \Delta E_{coop}$	$\approx 0 - 10$	5

The left column shows all nondimensional parameters that appear in the final equations. In the middle are their definition and estimated values. They are based on $\gamma_p = 5 \cdot 10^{-4} \text{ s}^{-1}$ and $K_a = 5 \cdot 10^{-5} \text{ M}$. Note that the range of the three MWC parameters depends on each other, but they can still be chosen independently. The range given for $[c]_a$ denotes the estimated “interesting” range in which switching happens, but $[c]_a$ can of course exceed these values. Details on the parameters and their estimation can be found in B.1. Finally, the last column shows the value that we use for the rest of this paper, unless otherwise noted. An explanation of this choice will follow in the next section.

3.5 Results and discussion

In the modeling process in the previous section, we have obtained three coupled differential equations. In this section, we will analyze these equations with deterministic methods and stochastic simulations. Analytical closed-form solutions could not be obtained and would, if they existed, probably not be helpful due to their large complexity. Finding such solutions requires solving a fifth order algebraic equation.

Deterministic phase portraits

A standard way to analyze dynamical systems deterministically is to plot phase portraits. In the following, we present such plots where the state variables are the mRNA, the protein, and the xanthosine concentration. Note that for the low copy numbers that can occur in our system, a deterministic analysis is not necessarily meaningful. Nevertheless, in our case we find that stochastic simulations agree well with the deterministic results. Thus, deterministic phase portraits are a valid starting point.

From a 3D to a 2D system. Figure 3.5A shows the 3D phase portrait for a representative set of parameters (shown in Table 3.1), whose choice is explained below. The plot looks rather complicated at first but can be understood intuitively. The three surfaces are the nullcline surfaces and the gray lines point in the direction in which the dynamical system moves at each point. The surfaces intersect in three points, which are the steady-state solutions of the dynamical system. For this choice of parameters, the system first flows towards the mRNA nullcline (independent of the initial condition), then it moves along that surface to the intersection with the protein nullcline, and lastly, it moves along that intersection line to one of the three intersection points of all three surfaces.

It is important to point out that, for a different set of parameters, the dynamics can be quite different. There are, for example, scenarios where the xanthosine kinetics are roughly as fast as the mRNA kinetics and the dynamics unfolds in two steps: first to the intersection of the mRNA and the xanthosine nullcline, then along that curve to the protein nullcline and thereby to a fixed point.

A usual simplification with genetic circuits like this is to assume the mRNA concentration to be at steady-state, i.e., to write $\frac{d[m]_a}{dt} = 0$ and solve this for $[m]_a([p]_a, [x]_a)$ to simplify the 3D to a 2D system. This restricts the dynamics to the green surface in our plot, which is reasonable here because as explained above, the system first

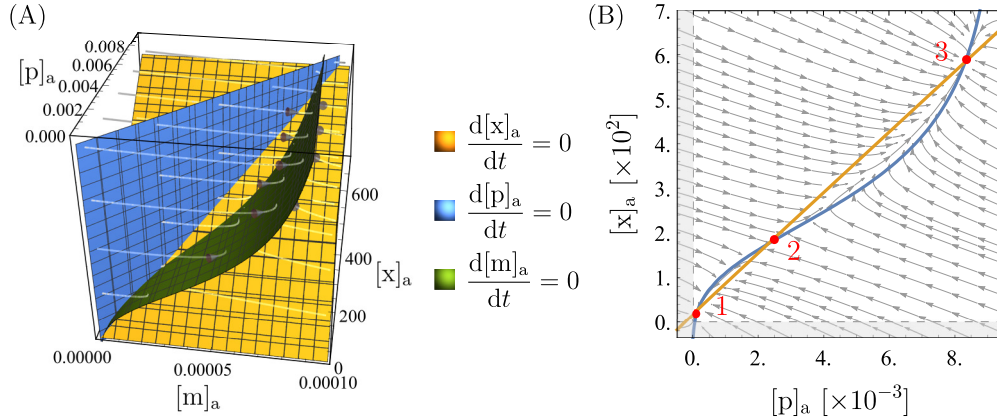


Figure 3.5: Phase portraits showing bistability. 3D and 2D phase portraits for one set of parameters that leads to bistability. The parameter values are listed in Table 3.1. Note that all the concentrations ($[m]_a$, $[p]_a$, $[x]_a$) are measured in units of $K_a = 5 \cdot 10^4$ nM. The surfaces in (A) and the curves in (B) are the nullclines of the state variables, and their intersection points, marked in red in (B), are the steady-state solutions of the system. The region shaded in gray in (B) leads to negative concentrations and is unphysical. A vector plot of (B) that also shows the magnitude of flow at each point can be found in B.1.

flows towards that surface before either the protein or the xanthosine concentration changes significantly. However, as already pointed out, this is different for other parameter values, and thus, this assumption does not hold in general. If the xanthosine dynamics are faster than the mRNA dynamics, the system first flows towards the xanthosine nullcline. In that case, forcing it onto the mRNA nullcline leads to significant changes in the dynamics.

Nevertheless, the steady-state solutions and the qualitative features that we address in this paper remain the same. Because the 3D plots are rather hard to read, we will, in the following, make the compromise to show a 2D version of the phase portraits but ensure that all of our statements also hold true in 3D space. As explained above, it makes the most sense here to do this by setting $\frac{d[m]_a}{d\tau} = 0$. The resulting equations can be found in B.1. In particular, we define $\rho := \frac{\rho_m \rho_p}{\gamma_{mp}}$ for everything that follows.

Bistability. We map the mRNA nullcline surface (green in Figure 3.5A) onto a plane to show it as the 2D plot in Figure 3.5B. From this 2D plot, it can clearly be seen that for the chosen parameters, there are three steady-state solutions. Because the system is restricted to the mRNA nullcline surface, these steady-state solutions are the same as those in the 3D plot ($\frac{d[m]_a}{d\tau} = 0$ on the nullcline and $\frac{d[p]_a}{d\tau} = 0$, $\frac{d[x]_a}{d\tau} = 0$ for the 2D fixed points). One can see from the vector field that the two

outer fixed points (labeled 1 and 3) are stable and the middle one (labeled 2) is unstable and serves as a sort of “switch-point” between the other two. This means that there are two stable states the cell can be in, one at high (point 3) and one at low (point 1) expression. As a result, there is bistability and the distribution of expression among cells can be bimodal, depending on initial conditions.

The bistability corresponds to the experimental observations, so the model passes this sanity check. Furthermore, the xanthosine and protein concentrations at the upper fixed point have the expected order of magnitude: the xanthosine concentration is roughly 10 – 100 mM, and there are roughly 500 proteins, which is just a bit lower than what was measured for the number of Nup transporters [16] which fulfill a similar purpose. We do not have well founded expectations for the other fixed points, so no comparison can be made here. Nevertheless, the orders of magnitude at the lower fixed point – roughly 1 – 10 nM of xanthosine and around 5 proteins – seem quite reasonable. Note that $[x]_a \approx [c]_a$ at the lower fixed point because there is only weak accumulation due to Nup and a few XapB transporters.

As already mentioned, we are working with one specific set of parameters here and we will now explain this choice of values. Firstly, they were picked roughly in the middle of the range that was estimated beforehand for this parameter (see Table 3.1 and B.1). Secondly, we chose parameters that allow clear bistability in the phase portraits as well as in the stochastic simulations (see later), which, of course, is not the case for any possible choice of parameters. Thirdly, by the corresponding choice of parameters it was ensured that the mRNA number per cell at the “switch-point” is around 1: this is large enough to enable the system to clearly resolve the two stable fixed points (as we will see from the stochastic simulations later on), but is low enough to lead to mean mRNA numbers that are very reasonable (of order 10^{-2} to 10^1 transcripts per gene per cell in *E. coli* [28]). The protein and xanthosine concentrations followed from this, but with some variation in the parameters they could still be tuned to a certain extent.

We point out that we have not observed any oscillations in the system. Intuitively, they might be expected when the XapA rate is significantly larger than the XapB rate, but it turns out that oscillations cannot be obtained. Why they do not occur can be understood when looking at the regions that are bounded by all three nullclines: on these boundaries, the streamlines point into the bounded regions, so deterministically, they serve as trapping regions from which the system cannot escape. Once inside, the only possible trajectory is non-oscillatory flow towards the stable fixed

point.

For a different set of parameters, the orders of magnitude in the plots and even the qualitative behavior can change. In the following, we will discuss some interesting features of the system that can be observed through the phase portraits.

The extracellular xanthosine concentration. The parameter that is the experimentally most easily tunable and biologically the most relevant is the extracellular xanthosine concentration. When it is increased in experiments, the cells go from (1.) all being in the low expression state to (2.) the population being in a mixed state with some cells in a low expression state and others in a high expression state (all-or-none phenomenon) and then to (3.) all being in the high expression state (see section “experimental motivation”). If our model is correct, it should exhibit the same qualitative behavior. Indeed we find exactly this: as can be seen in Fig 3.6, increasing $[c]_a$ makes the high stable fixed point appear and then, for even higher $[c]_a$, the lower one disappears. Thus, for low $[c]_a$ the only stable point of the system is at low expression, and for high $[c]_a$ there is only high expression. In between, there are two deterministically stable expression levels. Another means of visualizing this is with an bifurcation diagram, which we discuss in detail later.

Furthermore, we found that in the absence of xanthosine, i.e., setting $[c]_a = 0$ (not shown here), there are roughly 2-3 copies of XapA and XapB, which agrees very well with measurements, where around 2 copies per cell were found [16]. In addition, the parameter $K_{\chi A}$ (dissociation constant of xanthosine from active XapR) can be tuned such that the extracellular xanthosine concentration $[c]_a$ in the switching-regime is similar to that in the experiment. It was found that the cell only adapts at very high xanthosine concentrations of almost a millimolar [22] which is not completely unexpected when recalling that for *lac*, cells also limit themselves to glucose as long as possible. Interestingly, because there is no parameter other than $K_{\chi A}$ that tunes the critical value of $[c]_a$, this tells us that $K_{\chi A}$ is large as argued in the estimation of $K_{\chi A}$ in B.1. Thus, we predict that the interaction between xanthosine and XapR should be weak.

The roles of XapA and XapB. While it is clear that the bistability in the model system is due to the feedback loop from XapA and XapB, it is not intuitively clear if both XapA and XapB are necessary. The model implies that the bistability is due to XapB only. XapA is neither sufficient nor necessary and, within the estimated parameter regime, does not even have a significant influence on the system. This

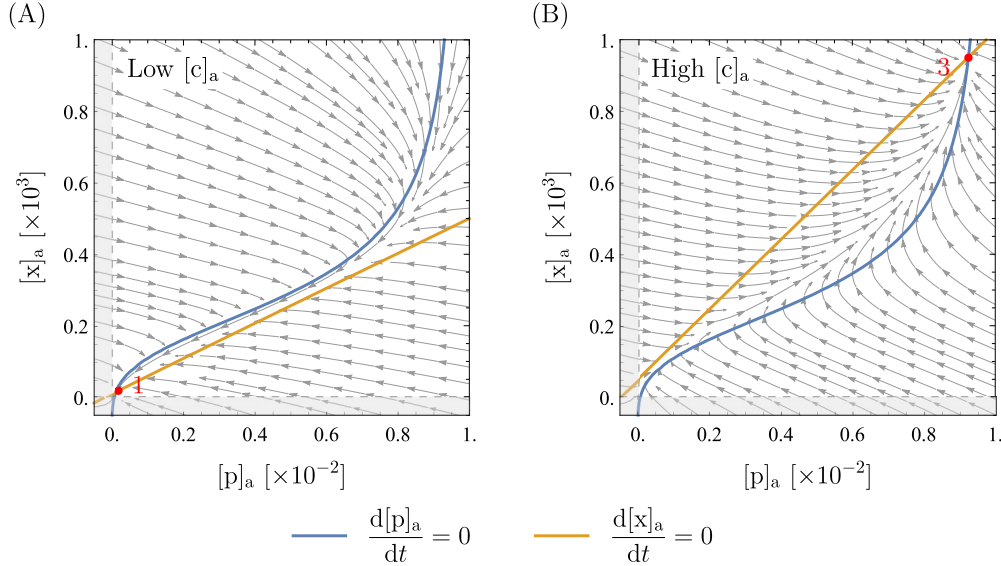


Figure 3.6: **Phase portraits for different extracellular xanthosine concentrations.** All parameters but $[c]_a$ are as presented in Table 3.1. The extracellular xanthosine concentration in these plots is $[c]_a = 7$ in (A) and $[c]_a = 40$ in (B) (recall that $[c]_a := \frac{c}{K_a}$ with $K_a = 5 \cdot 10^{-5}$ M, so $[c]_a$ is dimensionless). Tuning $[c]_a$ moves the orange line (xanthosine nullcline), but the blue curve (mRNA nullcline) is unchanged (see also B.1). It can clearly be seen that in (A) there is only the lower fixed point (fixed point number 1), whereas in (B) there is only the upper one (fixed point number 3). In between lies the bistable case that was shown in Fig 3.5.

can be seen from the plots in Fig 3.7. Degradation of xanthosine by XapA lowers the xanthosine and protein concentration at the upper fixed point by a small amount and could, in principle, thereby make the high-expression solution vanish. For our choice of all other parameters, bistability only vanishes for $k_\alpha > 10^4$ which is far from what has been measured. However, a higher effective rate could, in principle, be achieved by different translation rates of XapA and XapB (see simplifications of the model in B.1). Hence, we cannot exclude the possibility that XapA becomes so strong that it makes bistability impossible, but this is an extreme case. XapB, on the other hand, is essential; without it the system only has the one fixed point at low expression.²

For a cell, the minimal effect of XapA on bistability is a useful feature: by coupling

² Seeger et. al.[20] observed that $\Delta xapB$ mutants could survive, but grew extremely slowly, with xanthosine as the only carbon source. This suggests that low-affinity import of xanthosine by NupC and NupG is sufficient to sustain slow growth, but insufficient to serve as a stand-in for XapB. Fig 3.7 supports this supposition: in our model, with $xapB$ removed, the switch never activates and the cells are forced to survive with an extremely meager quantity of XapA to metabolize the abundant xanthosine.

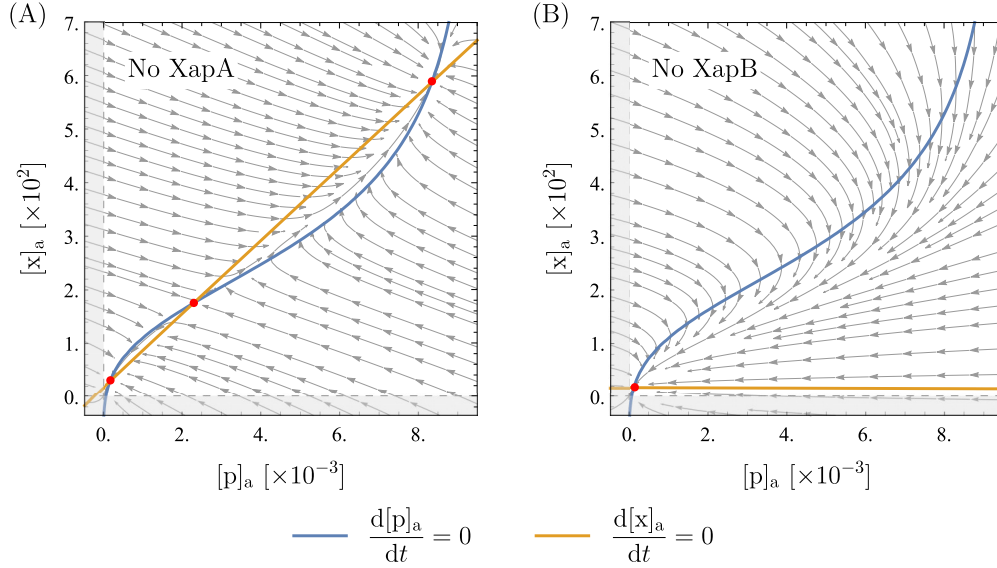


Figure 3.7: Phase portraits without XapA or XapB. All parameters are as presented in Table 3.1. In (A), the XapA term was removed from the kinetic equations. In (B), the equations lack the two terms from XapB. These plots clearly show that XapA has almost no influence on the qualitative behavior of the system (i.e. bistability and the order of magnitudes), but XapB is the essential feature for bistability.

XapA and XapB on an operon, XapA is switched on and off together with XapB but it does not significantly disturb this adaptation mechanism, while its kinetic parameters and expression levels can be chosen somewhat freely as necessary for metabolism. By having a membrane transporter gene on an operon whose expression is activated by the transporter substrate, the expression of a whole set of enzymes can be turned on and off depending on the presence of the substrate. It seems likely that this mechanism of short-term adaptation of a single cell to its environment may be used by cells for many metabolic processes.

The role of cooperativity. The model has two (putatively) cooperatively interacting binding sites for XapR on the *xapAB* promoter and two cooperative binding sites on XapR for xanthosine. It is interesting to consider whether the cooperativity is a necessary feature for bistability. This question is motivated by the importance of cooperativity in “typical” genetic switches[2, 29].

If, as a purely theoretical consideration, we remove either the second xanthosine binding site on XapR or the second XapR binding site on the promoter, leaving cooperativity in only one component of the system, we find that the system still has a bistable parameter regime. However, this bistable parameter range is smaller than in

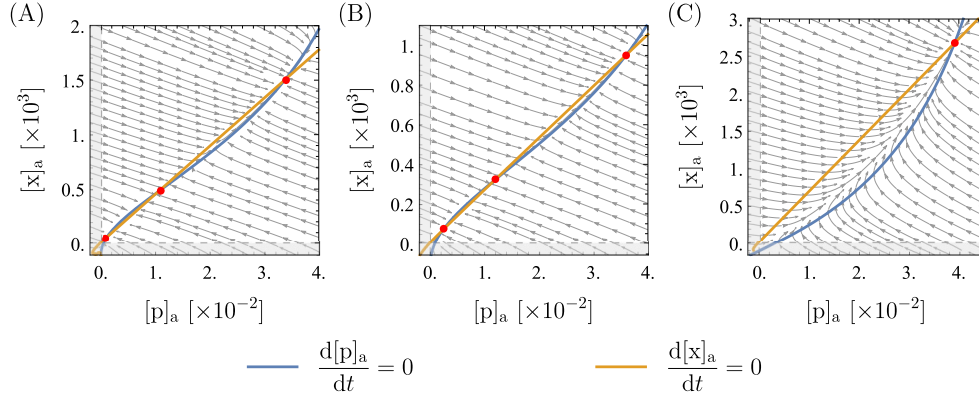


Figure 3.8: Phase portraits for less or no cooperativity. Most parameters are as presented in Table 3.1, changes are mentioned below. Fixed points are marked in red. In (A), there is only one xanthosine binding site on XapR and everything unchanged for the XapR-promoter binding. Two parameters are changed: $\rho = 0.07$ and $[c]_a = 6$. This is necessary to compensate for the weaker induction such that the system is bistable. In (B), there is only one XapR binding site on the promoter and everything is unchanged for the xanthosine-XapR binding. Two parameters are changed: $\rho = 0.13$ and $[c]_a = 3$. In (C), there is only one xanthosine binding site on XapR and also only one XapR binding site on the the promoter. Two parameters are changed: $\rho = 0.1$ and $[XapR]_R = 5$. Whereas bistability is retained in (A) and (B), it cannot be obtained anymore in (C).

the original model, which makes the system less stable: small stochastic fluctuations in the parameter values can collapse the system to monostability, possibly leaving it in the wrong state and without its ability to adapt. But only when the second binding site is removed in both places, leaving no cooperativity in the system, do we find that it is insufficiently non-linear to produce bistability. An example of the three scenarios (only cooperative XapR, only cooperative promoter, no cooperativity) can be seen in Fig 3.8. It follows that there need to be either two xanthosine binding sites on XapR or two XapR binding sites on the promoter (or both) in order to obtain a switch-like behavior.

One can also ask how much cooperative interaction is needed between the two binding sites. For the promoter, the amount of cooperativity is given by ΔE_{coop} in our model, and we find that setting $\Delta E_{\text{coop}} = 0$ has almost no influence on the phase diagrams. For XapR, we cannot test how much interaction is needed: the two binding sites interact indirectly, because the active state is much likelier if two xanthosine molecules are bound, and thus there is no continuous tuning parameter for the cooperative interaction like ΔE_{coop} in the case of the promoter.

Note that we are not writing Hill equations and measuring cooperativity in terms of the Hill coefficient. If Hill equations were to be used for the modeling, the Hill coefficient could have values between 1 and 2, which would yield bistability for large enough values, but not for lower ones. This could be investigated more rigorously similar to the analysis of a simple genetic switch in [29]. However, we refrain from looking for a minimal Hill coefficient in our system, because we do not find this very insightful. Hill equations only describe some specific limit cases of cooperative systems, but for example do not account for interaction energies and assume the partially bound states (e.g. only one XapR bound to the promoter) to never be populated. We suggest that cooperativity should be explored more in-depth and a more rigorous analysis of the role of cooperativity in simple genetic switches should be done before returning to more complex systems like this one.

Stochastic simulations

Stochastic simulations of the full 3-dimensional system of mRNA, protein, and xanthosine were run for comparison with the deterministic results. In B.1, we present the underlying chemical master equation of the system. Because of the two different fixed points at low and high expression, the protein copy numbers in the problem vary from less than five to several thousand. Even worse, xanthosine copy numbers may range as high as 10^7 at the high expression fixed point. For such large copy numbers, the number of reaction firings that must be simulated with Gillespie's classical algorithm [30] leads to an impractical computational cost. This would make Gillespie's τ -leap algorithm [31] ideal for the high expression state. On the other hand, τ -leaping cannot be used for the small protein copy numbers in the low expression state, or the mRNA copy number which remains of order ten or less in both states. For these reasons, we chose to work with the algorithm described in [32], a hybrid form between Gillespie's classical and his τ -leap algorithm. We gratefully worked with the Python implementation of this algorithm in *StochPy*, version 2.3 [33].

Note that we neglect stochastic fluctuations in any of our parameters, in particular in the XapR copy number. The latter is on the order of 10, but because of the long life span and rare expression of proteins like XapR, we expect the influence of the simplification on our results to be rather small. Nevertheless, the overall stochastic fluctuations are expected to be larger in the real system than in our simulations.

Bimodality and the extracellular xanthosine concentration. The stochastic approach results in the same bimodal distributions that were already seen in the deterministic investigation and experimental studies. Fig 3.9 shows the distribution of protein expression found in the simulations for different values of the extracellular xanthosine concentration. The parameters that were used are the same as in the previous section (listed in Table 3.1). However, we also find good agreement of the simulations and the deterministic results for other sets of parameters. To obtain the distributions, we ran the simulation 5000 times for a simulated time of 10^6 s each and started at an mRNA, protein and intracellular xanthosine count of 0.

The results agree very well with the deterministic fixed points and experiments: the mean numbers of mRNA, protein, and xanthosine in the stochastic results are as predicted from the phase portraits. It does, however, become clear that the phase portraits do not tell whether the cells will actually populate both the high and the low expression state, because they do not show the effective barrier height between the two states. In Figure 3.9(A), a deterministically bistable scenario is shown where the cells never switched to the high expression state during the run time of our simulations. This result implies that bistability, characterized by two separate *stable* steady-states, largely persists even in the presence of stochastic fluctuations, suggesting that the deterministic picture is remarkably effective. In other words, the fact that bimodality only occurs with some fine tuning of parameters means that the circuit is a strong switch and the deterministic picture is sufficient except in a small region of parameter space. Whether or not the system is actually in that region of parameter space remains a question for future experiments.

To elaborate on the preceding statements, we found that the two lower fixed points (marked as 1 and 2 in Fig 3.5) need to be very close like in Figure 3.9(B) to give bimodality. For lower $[c]_a$, meaning larger difference between the concentrations at the first and second fixed point, almost no switching was observed. Of course, switching is also a matter of the waiting time and stochastic effects: if one waits for long enough, it should eventually occur. However, switching times of more than several hours are not at the center of this investigation and would mean that switching is extremely unlikely. There are two aspects that become relevant in this context that we neglect in our analysis but briefly mention here: transcription and translation bursts lead to higher stochasticity and cell division leads to some discontinuity in the process. The effect of bursts is addressed in B.1.

Note that while the deterministic analysis assumes the variables to be continuous,

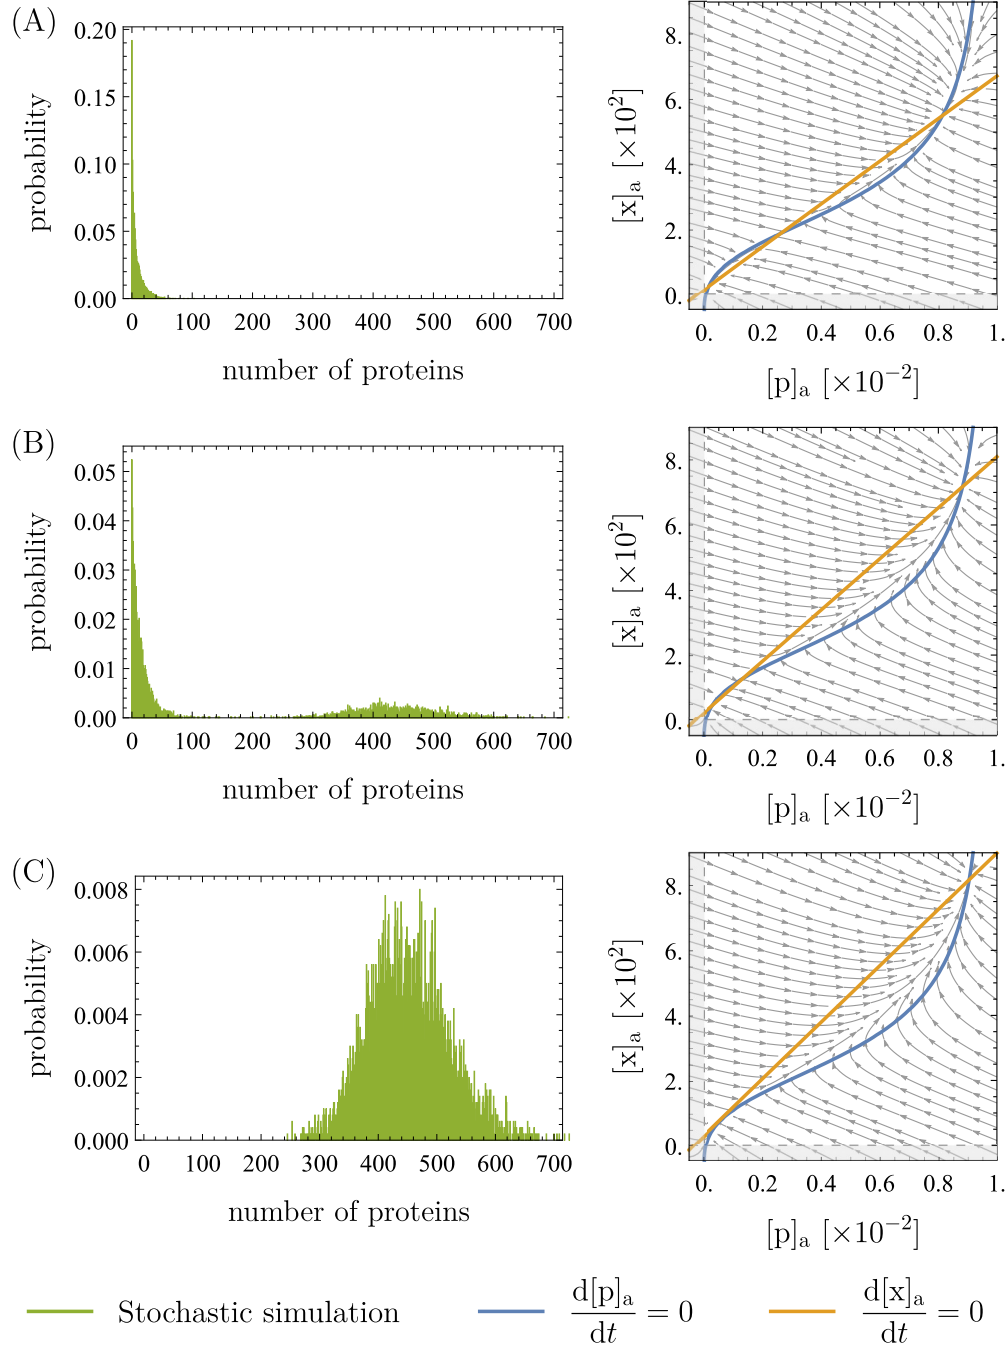


Figure 3.9: Distributions from stochastic simulations and the corresponding deterministic phase portraits. Apart from $[c]_a$, the parameters are the same as in Table 3.1. The phase portraits were obtained from the deterministic system similarly to those presented in the previous sections. For the distributions, the simulations were run 5000 times for 10^6 s each (simulated time) and started at a mRNA, protein and intracellular xanthosine count of 0. We show the two cases of unimodality (low expression in (A) and high expression in (C)) as well as the case of bimodality in (B). The values of $[c]_a$ are 12 in (A), 18.5 in (B), and 25 in (C) (recall that $[c]_a := \frac{c}{K_a}$ with $K_a = 5 \cdot 10^{-5}$ M, so $[c]_a$ is dimensionless). The output from the stochastic simulations is in good agreement with the concentrations at the fixed points in the deterministic phase portraits.

the simulations work with discrete numbers of mRNA, protein, and xanthosine. This per se is no problem, because the deterministic analysis describes the mean values and the simulation fluctuates around this mean. However, if the mRNA number of the third (high) fixed point is so low that stochastic fluctuations are larger than the difference in concentrations between the first and the second fixed point, the system may not be able to resolve the two points anymore. The tolerance to this is surprisingly large, though: In Fig 3.9, the distance between the first and the third fixed point is around 3 mRNA molecules. While this is around the size of the fluctuations in Fig 3.9, the latter become as large as 10 mRNA molecules when bursts are included in the simulation (see B.1) and yet, the system is able to resolve the fixed points very well.

Time evolution and adaptation times. Figure 3.10A shows the time evolution of the protein concentration for one typical run of the simulation. Again, the simulation was started with a mRNA, protein and intracellular xanthosine count of 0 and was run for a simulated time of $5 \cdot 10^5$ s. In this specific example, the fixed point was reached after roughly $1.5 \cdot 10^5$ s (adaptation time). This time varies: Figure 3.10B shows the distribution of adaptation times from 1000 runs of the simulation. In both figures, the trajectory obtained from integrating the deterministic system is shown for comparison. It clearly agrees well with the simulation on average.

Comparing this to experimentally expected timescales is difficult, because the adaptation time strongly depends on the extracellular xanthosine concentration. Experiments were always stopped after a few hours, and in this time, the cell population might not reach its steady-state expression distribution. Hence, the distribution could be bimodal when the experiment is stopped but become unimodal after further waiting. That way, extracellular xanthosine concentrations that are too high for deterministic bistability could lead to experimental bimodality if the experiment is stopped too early. In this case, the observed timescales would be shorter, which makes the comparison to our simulations even harder. Thus, we cannot say if it is problematic that the 10^5 s is larger than what was found in the experiment.

Nevertheless, we do warn the reader that the timescales in the simulations and even more so in the deterministic system should be taken with reservation. Fluctuations in the parameters are not considered here, and neither are cell divisions or the burstiness of transcription and translation. This means that stochasticity may be larger in the real system which should have an influence on the timescales and may shorten the

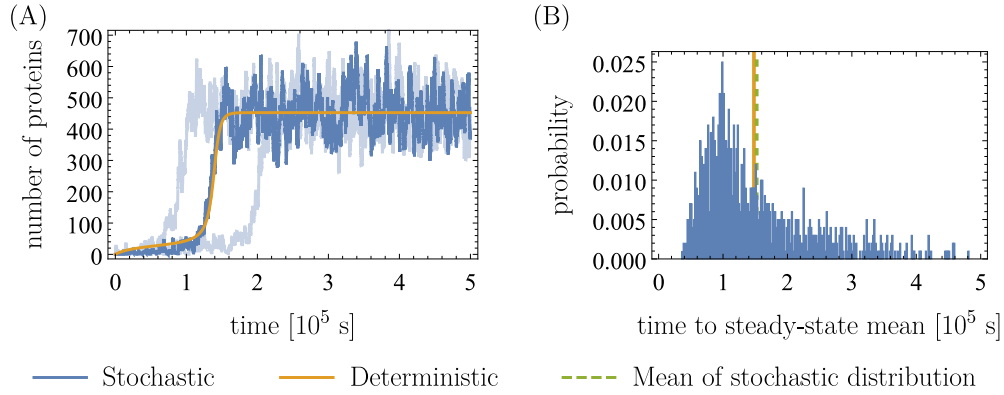


Figure 3.10: Stochastic and deterministic time evolution of protein (XapA/XapB) and adaptation time. (A) Shown in blue is the result from one typical run of the stochastic simulation, and in orange the trajectory obtained from solving the deterministic ODE's. In light blue, two more extreme runs of the simulation are shown for comparison. The simulation was run for $5 \cdot 10^5$ s each and started at an mRNA, protein and intracellular xanthosine count of 0. The parameters that were used are the same as in Table 3.1, the only exception being the extracellular xanthosine concentration, which was chosen to be $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M) just as in Figure 3.9(C). (B) In blue we plot the waiting time distribution for 1000 runs of the simulation (same conditions as (A)) to reach the steady-state mean. We define this time as the elapsed time when the protein copy number first reaches 90% of its value at the upper fixed point. To better visualize the bulk of the distribution, we excluded from the plot ~ 10 runs with adaptation times larger than $6 \cdot 10^5$ s. The green dashed line indicates the mean of the blue distribution. The orange line shows the corresponding deterministic time.

time until the fixed point is reached. However, including transcriptional bursts in the stochastic simulation changes little in the output: the qualitative behavior remains the same but fluctuations around the mean as well as in the adaptation time become larger and bimodality already occurs at lower $[c]_a$ (see B.1). Still, the fact that there are no big changes shows that the system is stable to stochastic perturbations and our particular assumptions should not be too significant.

Of course, the system moves to its mean steady-state much faster when the extracellular xanthosine concentration is further away from the bimodal regime (in analogy to a “critical point”). Note that it could well be that in reality, the extracellular xanthosine concentration is high enough to be in the regime where there only is the upper fixed point and thus no bistability or even bimodality. As the system reaches its mean steady-state faster in that regime, the bacteria could adapt more quickly.

Bifurcation diagram and hysteresis. In Figs 3.9 and 3.10, the simulations were started at initial intracellular concentrations of 0 to investigate what happens if xanthosine is suddenly added to the cell's environment. We can now ask the opposite question: what happens when xanthosine is removed from the extracellular environment? To answer this, the simulation was started with initially fully induced cells, i.e. at the mRNA, protein and intracellular xanthosine counts of the high fixed point in the corresponding phase portrait.

The distributions we obtained can be found in B.1. Here, we instead present the results in the form of a bifurcation diagram in Fig 3.11. In blue, the positions of the deterministic fixed points for the corresponding extracellular xanthosine concentration are shown. As explained before, there is only one fixed point for low and high values of $[c]_a$, but in between, there are three. In yellow and orange, the results from the stochastic simulation are shown, where the mean of each distribution was taken. For the orange points, the simulation was started at $[m]_a = [p]_a = [x]_a = 0$, leading to the distributions in Fig 3.9. In contrast, the yellow points result from starting the simulation at the $[m]_a$, $[p]_a$, and $[x]_a$ values of the high fixed point. Note that by taking the mean of each distribution, we hide the feature of bimodality. For that reason, we additionally show the approximate position of the two peaks in the bimodal distribution. We do this by estimating the position of the minimum between the two peaks of the distribution, splitting it into two parts there, and calculating the mean of each of the parts. Also note that it suffices to show the bifurcation diagram in one dimension (here protein concentration) because this uniquely determines the fixed points (recall from Fig 3.5, the fixed points are unique points in the 3D system).

The figure clearly shows the hysteresis in the system: there exist extracellular xanthosine concentrations where initially uninduced cells remain uninduced and initially induced cells remain induced. Only when the second and the third fixed point are very close can initially induced cells “switch off” to the uninduced state. This behavior is symmetric to the “switching on” in the previous paragraphs. The arrows in Fig 3.11 indicate the hysteresis loop.

In other words, cells only change their metabolism to xanthosine if enough of the latter is around, but after they have switched, this metabolic state is stable even if the xanthosine concentration decreases to a certain extent. This stability explains what was observed by Novick and Weiner [5] for the *lac* operon: when induced cells were transferred into lower concentrations of lactose, they remained induced, even though uninduced cells could not become induced at these concentrations.

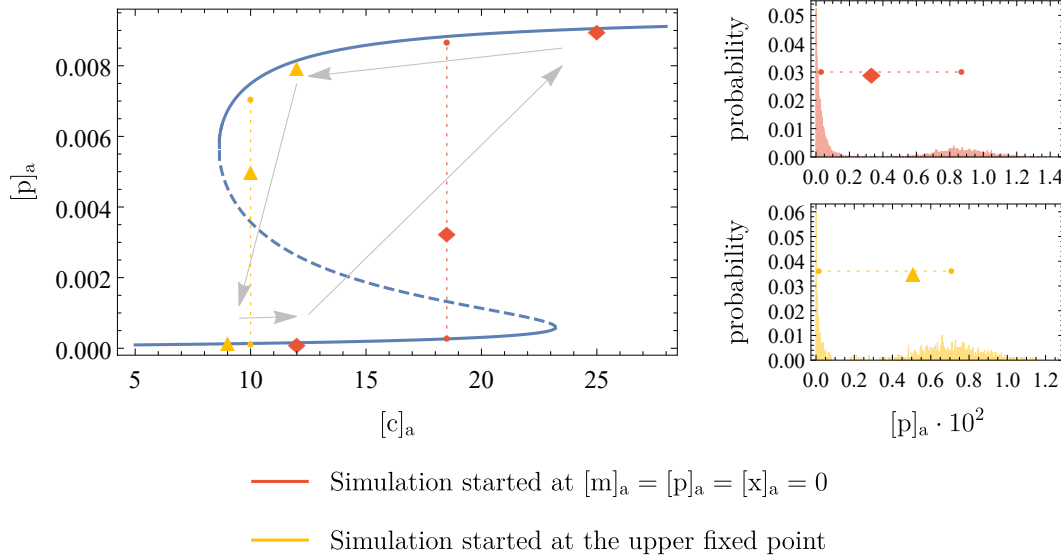


Figure 3.11: **Bifurcation diagram showing the hysteresis loop in the stochastic system.** The blue line shows the positions of the deterministic fixed points, while the dashed line indicates the instability of the middle one. The orange and yellow points show the mean of the stochastic simulations when started at zero and at the high fixed point, respectively. The positions of the two peaks in the bimodal distributions are indicated by smaller points, connected by dotted lines. To make this clearer, the bimodal distributions themselves are shown on the right. The arrows illustrate the hysteresis loop in the stochastic simulations.

In addition to the hysteresis, Fig 3.11 also illustrates the astonishingly good agreement of the stochastic simulation and the deterministic results, despite having copy numbers below 10 in some cases. Although one should be cautious about this because of higher stochasticity in the real biological system (addressed above), the result does show that the switch-like feature in the circuit is strong and stable.

3.6 Conclusion

Here, we propose a simple model for genetic circuits containing a membrane transporter whose gene expression is, directly or indirectly, activated by its substrate. We have shown that such a system can be bistable and thus work as a genetic switch which reacts to the extracellular concentration of the relevant metabolite. This switch has very useful biological features. First, coupling of the transporter with, for example, an enzyme which metabolizes the transporter substrate creates a genetic switch that enables short-term adaptation of the cell's metabolism to its environment. Second, the switch is stabilized by hysteresis effects when the extracellular substrate concentration decreases, which explains previous experimental

findings. Furthermore, our simulations show that the switch-like behavior is very robust.

However, we have found that no bistability can emerge from the genetic circuit unless at least one component has two binding sites for its activator. Additional binding sites or cooperativity seem to increase the stability of the switch. In addition, simply knowing the experimental switching concentration of xanthosine permits us, for example, to infer the approximate value of the dissociation constant between the transcription factor XapR and the inducer xanthosine. The value we infer is roughly one to two orders of magnitude larger than what has been measured for LacI and IPTG [23], meaning the interaction of XapR and xanthosine is rather weak.

Phase diagrams, showing for which parameters the system is bistable and for which there is only the lower or the upper stable fixed point, could be calculated from arguments made in [29]. However, the simulations showed that in the deterministically bistable regime, which fixed point(s) the system occupies is dependent on initial conditions, which is why we have refrained from showing such phase diagrams. Furthermore, the timescales in the problem could be investigated more thoroughly, for example the dependence of the switching time on $[c]_a$, but such an analysis should probably account for cell divisions and fluctuations in the parameters, which is not straightforward. Lastly, it could be interesting to investigate the magnitude of the fluctuations around the fixed point away from and near the bifurcations in $[c]_a$.

Despite the small copy numbers at the lower fixed point, the stochastic simulations are in excellent agreement with deterministic predictions. All model parameters could be reasonably estimated despite the paucity of experimental knowledge about the model system. The concentrations of mRNA, protein, and xanthosine at the fixed points as well as all qualitative features are as expected from similar systems and the few experiments on the *xap* circuit. These points suggest that the model captures the relevant components of the system correctly and is able to describe its dynamics. The modeling results let us, to some extent, understand how the biological system can achieve its function. By keeping the model as minimal as possible, but still modeling every part explicitly with an appropriate complexity, we can investigate the interesting features while still being able to understand the influence of all parameters and their interplay intuitively.

With the framework given in this text, it should be straightforward to model other promoters, regulatory pathways or enzymes and thereby adapt the model to other genes and metabolites. Examples include *lac*, *ara*, and *xyl*, but we suspect that

many if not most metabolic processes involve the adaptation mechanism that we have investigated here, and that much can be understood about them through our model. This apparent success demonstrates once more that even for broadly unknown systems, rigorous physical modeling can potentially offer an efficient way to gain a very thorough understanding of the behavior of the system.

Materials and Methods

Bacterial Strains. Strains used in this work derive from *E. coli* K12 MG1655 with the *lac* operon removed and were constructed similarly to those used in [23, 34]. The $\Delta xapABR$ knockout was generated using the approach of Datsenko and Wanner [35]. For the fluorescent reporter constructs the *xapAB* promoter region, with one or both XapR binding sites, were extracted from *E. coli* K12 MG1655 by PCR. These were cloned into the YFP-expressing reporter plasmid pZS25 [36], containing kanamycin resistance, and integrated at the *galK* locus using λ Red recombineering [37]. A xapR-mCherry fusion was constructed in a pZS31 plasmid and was integrated into the *ycbN* locus, also using λ Red recombineering. Expression of the xapR-mCherry fusion was regulated by TetR expressed extrachromosomally from a pZS3-PN25 plasmid. Expression of the xapR-mCherry fusion was induced by the addition of 10 ng/mL anhydro-tetracycline. A similar expression system was used in [38] More details can be found in B.1.

Flow Cytometry Measurements. Experimental measurements in Fig. 2 were obtained via flow cytometry and were performed in a similar manner as described in [23]. Briefly, wild-type, $\Delta xapABR$, and $\Delta xapAB$ were grown to saturation in rich LB Miller (BD Medical) medium and were diluted 1:1000 into M9 minimal medium supplemented with 0.5% (w/v) glucose and the appropriate concentration of xanthosine. Cultures were allowed to grow at 37 °C for six to eight hours to an $OD_{600\text{ nm}}$ of ≈ 0.3 . At this point, the cultures were diluted 1:10 into M9 + 0.5% glucose and were immediately analyzed via a MacsQuant VYB Flow Cytometry (MiltenyiBiotec). Measurement files were exported to CSV file formats and analyzed as previously described in [23].

Supporting information

S1 Text. The aforementioned further information. Discussion of simplifications in the model, parameter estimation, elaborations on the results, and the chemical master equation of this circuit. Experimental materials and methods.

Acknowledgments

We thank Jane Kondev and Jin Wang for interesting discussions, Patrick Lenggenhager for his help with *Mathematica* issues, and Nigel Orme for assistance with illustrations. Portions of the experiments reported here were carried out at the Physiology Course at the Marine Biological Laboratory in Woods Hole, MA, operated by the University of Chicago. This work was supported by 1R35 GM118043-01 Maximizing Investigators' Research Award (MIRA) (to R.P.), and the Werner Siemens Foundation through the Swiss Study Foundation (K.S.L.). This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1745301 (M.J.M.).

References

- ¹F. Jacob and J. Monod, "Genetic regulatory mechanisms in the synthesis of proteins", *Journal of Molecular Biology* **3**, 318–356 (1961).
- ²T. S. Gardner, C. R. Cantor, and J. J. Collins, "Construction of a genetic toggle switch in *Escherichia coli*", *Nature* **403**, 339–342 (2000).
- ³D. M. Wolf and F. H. Eeckman, "On the Relationship between Genomic Regulatory Element Organization and Gene Regulatory Dynamics", *Journal of Theoretical Biology* **195**, 167–186 (1998).
- ⁴P. Wong, S. Gladney, and J. Keasling, "Mathematical Model of the lac Operon: Inducer Exclusion, Catabolite Repression, and Diauxic Growth on Glucose and Lactose", *Biotechnology Progress* **13**, 132–143 (1997).
- ⁵A. Novick and M. Weiner, "ENZYME INDUCTION AS AN ALL-OR-NONE PHENOMENON", *Proceedings of the National Academy of Sciences* **43**, 553–566 (1957).
- ⁶M. Santillán and M. C. Mackey, "Quantitative approaches to the study of bistability in the *lac* operon of *Escherichia coli*", *Journal of The Royal Society Interface* **5**, 10.1098/rsif.2008.0086.focus (2008).
- ⁷E. M. Ozbudak, M. Thattai, H. N. Lim, B. I. Shraiman, and A. van Oudenaarden, "Multistability in the lactose utilization network of *Escherichia coli*", *Nature* **427**, 737–740 (2004).
- ⁸A. Narang and S. S. Pilyugin, "Bistability of the lac Operon During Growth of *Escherichia coli* on Lactose and Lactose + Glucose", *Bulletin of Mathematical Biology* **70**, 1032–1064 (2008).
- ⁹P. J. Choi, L. Cai, K. Frieda, and X. S. Xie, "A Stochastic Single-Molecule Event Triggers Phenotype Switching of a Bacterial Cell", *Science* **322**, 442–446 (2008).

- ¹⁰G. Fritz, J. A. Megerle, S. A. Westermayer, D. Brick, R. Heermann, K. Jung, J. O. Rädler, and U. Gerland, “Single Cell Kinetics of Phenotypic Switching in the Arabinose Utilization System of *E. coli*”, *PLoS ONE* **9**, edited by C. Herman, e89532 (2014).
- ¹¹A. Jenkins and M. Macauley, “Bistability and Asynchrony in a Boolean Model of the l-arabinose Operon in *Escherichia coli*”, *Bulletin of Mathematical Biology* **79**, 1778–1795 (2017).
- ¹²D. A. Siegle and J. C. Hu, “Gene expression from plasmids containing the araBAD promoter at subsaturating inducer concentrations represents mixed populations”, *Proceedings of the National Academy of Sciences* **94**, 8168–8172 (1997).
- ¹³J.-S. Bae, T.-H. Kim, M.-Y. Kim, J.-M. Park, and Y.-H. Ahn, “Transcriptional Regulation of Glucose Sensors in Pancreatic β -Cells and Liver: An Update”, *Sensors* **10**, 5031–5053 (2010).
- ¹⁴M. Tiedge and S. Lenzen, “Regulation of glucokinase and GLUT-2 glucose-transporter gene expression in pancreatic B-cells”, *Biochemical Journal* **279**, 899–901 (1991).
- ¹⁵A. Schmidt et al., “The quantitative and condition-dependent *Escherichia coli* proteome”, *Nature Biotechnology* **34**, 104–110 (2016).
- ¹⁶G.-W. Li, D. Burkhardt, C. Gross, and J. S. Weissman, “Quantifying Absolute Protein Synthesis Rates Reveals Principles Underlying Allocation of Cellular Resources”, *Cell* **157**, 624–635 (2014).
- ¹⁷S. Song and C. Park, “Organization and regulation of the D-xylose operons in *Escherichia coli* K-12: XylR acts as a transcriptional activator.”, *Journal of bacteriology* **179**, 7025–7032 (1997).
- ¹⁸R. S. Buxton, K. Hammer-Jespersen, and P. Valentin-Hansen, “A second purine nucleoside phosphorylase in *Escherichia coli* K-12: I. Xanthosine phosphorylase regulatory mutants isolated as secondary-site revertants of a deoD mutant”, *Molecular and General Genetics MGG* **179**, 331–340 (1980).
- ¹⁹K. Hammer-Jespersen, R. S. Buxton, and T. D. Hansen, “A second purine nucleoside phosphorylase in *Escherichia coli* K-12: II. Properties of xanthosine phosphorylase and its induction by xanthosine”, *Molecular and General Genetics MGG* **179**, 341–348 (1980).
- ²⁰C. Seeger, C. Poulsen, and G. Dandanell, “Identification and characterization of genes (xapA, xapB, and xapR) involved in xanthosine catabolism in *Escherichia coli*.”, *Journal of bacteriology* **177**, 5506–5516 (1995).
- ²¹C. Jørgensen and G. Dandanell, “Isolation and characterization of mutations in the *Escherichia coli* regulatory protein XapR”, *Journal of Bacteriology* **181**, 4397–4403 (1999).

- ²²M. H. H. Nørholm and G. Dandanell, “Specificity and Topology of the Escherichia coli Xanthosine Permease, a Representative of the NHS Subfamily of the Major Facilitator Superfamily”, *Journal of Bacteriology* **183**, 4900–4904 (2001).
- ²³M. Razo-Mejia, S. L. Barnes, N. M. Belliveau, G. Chure, T. Einav, M. Lewis, and R. Phillips, “Tuning Transcriptional Regulation through Signaling: A Predictive Theory of Allosteric Induction”, *Cell Systems* **6**, 456–469.e10 (2018).
- ²⁴H. G. Garcia and R. Phillips, “Quantitative dissection of the simple repression input-output function”, *Proceedings of the National Academy of Sciences* **108**, 12173–12178 (2011).
- ²⁵G. Chure, M. Razo-Mejia, N. M. Belliveau, T. Einav, Z. A. Kaczmarek, S. L. Barnes, M. Lewis, and R. Phillips, “Predictive shifts in free energy couple mutations to their phenotypic consequences”, *Proceedings of the National Academy of Sciences* **116**, 18275–18284 (2019).
- ²⁶S. Marzen, H. G. Garcia, and R. Phillips, “Statistical Mechanics of Monod–Wyman–Changeux (MWC) Models”, *Journal of Molecular Biology* **425**, 1433–1460 (2013).
- ²⁷H. R. Kaback, “A chemiosmotic mechanism of symport”, *Proceedings of the National Academy of Sciences* **112**, 1259–1264 (2015).
- ²⁸R. Milo and R. Phillips, *Cell biology by the numbers* (Garland Science, Taylor & Francis Group, New York, NY, 2016), 356 pp.
- ²⁹J. L. Cherry and F. R. Adler, “How to make a Biological Switch”, *Journal of Theoretical Biology* **203**, 117–133 (2000).
- ³⁰D. T. Gillespie, “A general method for numerically simulating the stochastic time evolution of coupled chemical reactions”, *Journal of Computational Physics* **22**, 403–434 (1976).
- ³¹D. T. Gillespie, “Approximate accelerated stochastic simulation of chemically reacting systems”, *The Journal of Chemical Physics* **115**, 1716–1733 (2001).
- ³²Y. Cao, D. T. Gillespie, and L. R. Petzold, “Efficient step size selection for the tau-leaping simulation method”, *The Journal of Chemical Physics* **124**, 044109 (2006).
- ³³T. R. Maarleveld, B. G. Olivier, and F. J. Bruggeman, “StochPy: A Comprehensive, User-Friendly Tool for Simulating Stochastic Biological Processes”, *PLoS ONE* **8**, edited by J. Bourdon, e79345 (2013).
- ³⁴N. M. Belliveau, S. L. Barnes, W. T. Ireland, D. L. Jones, M. J. Sweredoski, A. Moradian, S. Hess, J. B. Kinney, and R. Phillips, “Systematic approach for dissecting the molecular mechanisms of transcriptional regulation in bacteria”, *Proceedings of the National Academy of Sciences of the United States of America* **115**, E4796–E4805 (2018).

- ³⁵K. A. Datsenko and B. L. Wanner, “One-step inactivation of chromosomal genes in *Escherichia coli* K-12 using PCR products”, *Proceedings of the National Academy of Sciences* **97**, 6640–6645 (2000).
- ³⁶R. Lutz and H. Bujard, “Independent and tight regulation of transcriptional units in *Escherichia coli* via the LacR/O, the TetR/O and AraC/I1-I2 regulatory elements”, *Nucleic Acids Research* **25**, 1203–1210 (1997).
- ³⁷S. K. Sharan, L. C. Thomason, S. G. Kuznetsov, and D. L. Court, “Recombineering: a homologous recombination-based method of genetic engineering”, *Nature Protocols* **4**, 206–223 (2009).
- ³⁸R. C. Brewster, F. M. Weinert, H. G. Garcia, D. Song, M. Rydenfelt, and R. Phillips, “The Transcription Factor Titration Effect Dictates Level of Gene Expression”, *Cell* **156**, 1312–1323 (2014).

Chapter 4

AN EFFICIENT REPRESENTATION FOR STATISTICAL MECHANICS OF MULTI-PARTICLE COMPLEXES IN BIOLOGY

4.1 Introduction

Motivation

Biology runs on multimeric complexes of proteins. Across all domains of life, less than half of proteins are found as monomers, with the rest existing as dimers or higher-order complexes [1]. And this accounting is merely of proteins' "default" configurations. It neglects the even larger transient complexes in which they participate, which are often their *raison d'être* - the preinitiation complexes for transcription and replication, spliceosomes, and GPCRs are prominent examples. In addition to protein-protein interactions, proteins' small-molecule binding partners are just as important to consider. Witness the power of allostery, as codified in, e.g., the MWC model [2, 3], where the binding of a regulatory small molecule at one site can change a protein's binding affinities or enzymatic rates at other active sites by orders of magnitude, effectively "transmuting" a protein from one species to another. From a theoretical view, the proliferation of possible biochemical complexes threatens to be overwhelming, and existing mathematical methods, at best, struggle to meet this challenge. How do we begin to model such complexity?

The chemical master equation is one natural starting point for modeling many biochemical systems, when our goal is a coarse-grained picture that neglects atomistic-level details. An exact solution specifies the probability distribution over available states of the system as a function of time, given some initial condition, but finding such a solution is rarely possible. Generating function methods can provide such a solution, for instance, as we used in Chapter 2, but generally apply only for limited classes of master equations [4]. In some cases, though an exact calculation of the full distribution may be impossible, an exact analytical calculation of the distribution's *moments* is possible (e.g., [5, 6]). And if merely specifying mean quantities of the steady-state distribution is sufficient, the King-Altman diagram method, independently discovered by Hill, intuitively solves this problem by representing the problem in graphical form [7, 8].

For systems with variable number of (classical) identical particles, one novel alternative formulation of a master equation uses Fock spaces, the vector spaces normally associated with quantum field theories, which offer a natural representation of the variable particle number. Several flavors of such Fock space formalisms have been developed. Doi offered the original formulation as a second-quantized description of reaction-diffusion systems [9, 10]. Several variants followed [11, 12], as well as a path integral description [13]. In particular, the methods of Grassberger and Scheunert [11] as well as Peliti [13] have been widely adopted [4, 14], especially in studies of diffusion-limited processes [15]. A sampling of more recent work has covered such diverse topics as population and epidemic models [16], predator-prey interactions [17], stem cell differentiation [18], and neural networks [19] with self-organized criticality [20] or transitions between metastable states [21]. Many such studies have taken advantage of the path-integral formulation to address the dynamic phase transitions, and their critical behavior, arising in such models.

However, for treating chemical systems that generate large multi-particle complexes, these formalisms rapidly become intractable. Many interesting chemical systems in biology and polymer physics are capable of generating vast (or even infinite) numbers of distinct complexes from a relatively small number of components and interaction rules. In principle, existing Fock space methods can be used to describe such systems but in practice, this is not feasible because existing formalisms treat each distinct multi-particle complex as its own species of particle, requiring manual enumeration of all possible complexes, along with their corresponding free energies, formation and decay rates, and so on. But often these complexes are built up from smaller molecules, which these methods fail to take advantage of, leading to enormous redundancy in merely defining the system. Some authors have undertaken the Herculean effort to enumerate lists of complexes, sometimes running to hundreds of entries [22, 23]. To our knowledge, there has been only one attempt to treat more than three species with the Doi-Peliti formalism, and the resulting field theory had so many fields that the only tractable approach was a computer simulation of the field theory, rather than the underlying master equations [18].

The intractable proliferation of complexes is a well-recognized problem in the context of molecular systems biology [24–26]. To address this issue computationally, formal grammars [27–31] and accompanying software [32–36] have been developed that enable “rule-based” simulations of biochemical systems. While such techniques are undoubtedly useful, we hold out hope for an intuitive “rule-based mathematics”

that allows one to work with such systems analytically, which has yet to be described.

Aside from the inelegance of enumerating all complexes, as the number of components grows there is a real danger that some complexes will be overlooked in the combinatorial explosion. While their neglect may or may not cripple the resulting theory, it is impossible to know *a priori* what the effect might be. Such situations demand a theory that can self-consistently account for all the necessary complexes. Such a formalism is the goal of the present work.

Sketch of the method

Our formalism introduces a Fock space similar in spirit though not in detail to Doi's [9] and Park and Park's [12]. Although our method is motivated these Fock space methods for nonequilibrium problems, and although a treatment of the full non-equilibrium problem would be the ultimate goal, we have found that even a simpler equilibrium formalism remains full of subtleties and surprises. Therefore, in this work we present an equilibrium treatment only and defer a non-equilibrium formalism for future work. The present work introduces two powerful innovations.

The first key is to model every particle in this formalism as existing in one of a large number of internal states. These internal states uniquely identify each particle and are essential for representing multi-particle complexes in terms of their components. These degrees of freedom are not in themselves of interest. They merely serve to enforce a classical notion of distinguishability while retaining the point particle idealization. Therefore, multiple molecules of the same type can exist at the same "point" in space, forming bound complexes with each other and with other species.

This leads us to the second key idea, which is to track binding sites on, interactions between, and conformational states of molecules separately from the molecules themselves. In other words, while one field creation operator "creates" the molecule, another field's creation operator might "occupy" a binding site on it, another might represent pairwise bonds between two such molecules, and yet another could transform its conformational state. This approach allows complexes to be assembled naturally by specifying intuitive "rules," rather than introducing whole complexes as entirely new species produced by reactions. For the present work, we will view monomers and complexes as zero-dimensional point-like idealizations, but this restriction can be relaxed. We believe occupied volume and steric interactions can be incorporated in the theory as another auxiliary field, though we also leave that for future work.

We emphasize two important features of our approach relative to existing work. First, previous Fock-space formalisms are only capable of creating and destroying particles, so complexes must be created “whole-cloth.” In contrast, our approach naturally and intuitively allows complexes to be assembled, disassembled, or have conformational states modified. Second, our approach offers a natural way to coarse-grain multi-particle complexes, as we will see in detail below.

Carrying out calculations directly with our operator formalism remains algebraically daunting for all but the simplest examples. To assist, we have developed a convenient diagrammatic approach that is equivalent, yet far more transparent. It should be noted that our diagram techniques, though obviously inspired by the familiar Feynman diagrams from quantum field theory, are not equivalent to them, unlike the diagrams arising from the Doi-Peliti approach. This will be made more clear below.

4.2 Formalism

Overview

Following [9], we equip our Fock space with an orthonormal basis \mathcal{S} of “pure” physical states. More will be said on the construction of \mathcal{S} below, but one can think of each state $s \in \mathcal{S}$ as an eigenstate in particle numbers, i.e., for each conceivable complex allowed in the model, s possesses a definite number of particles of that type. \mathcal{S} contains all possible such states s . The vector $|\psi\rangle$ that describes a system of interest is given by

$$|\psi\rangle = \sum_{s \in \mathcal{S}} p_s |s\rangle, \quad (4.1)$$

where p_s is the probability that the system will be found in state s . Every measurable quantity Q is represented by a corresponding operator \mathbb{Q} that is diagonal in \mathcal{S} , and the expectation value for this quantity is given in terms of $|\psi\rangle$ by $\langle Q \rangle = \langle \text{sum} | \mathbb{Q} | \psi \rangle$, where

$$|\text{sum}\rangle \equiv \sum_{s \in \mathcal{S}} |s\rangle \quad (4.2)$$

is referred to as the “sum vector.” $|\text{sum}\rangle$ can be thought of as an enumeration of all particle number eigenstates, or microstates in a grand canonical ensemble, spanning all possible particle numbers, of all possible complexes. Note that $\langle \text{sum} | \text{sum} \rangle \neq 1$, so unlike in quantum mechanics, there is no sense in which $|\text{sum}\rangle$ is normalized.

In thermal equilibrium, $|\psi\rangle$ is uniquely determined by $|\text{sum}\rangle$ and a “Hamiltonian”

operator \mathbb{H} that assigns a Gibbs free energy to each pure state:

$$|\psi\rangle = \frac{e^{-\mathbb{H}/kT}}{Z} |\text{sum}\rangle, \quad (4.3)$$

where the partition function Z is given by

$$Z \equiv \langle \text{sum} | e^{-\mathbb{H}/kT} | \text{sum} \rangle. \quad (4.4)$$

Much of our effort will be invested in properly constructing $|\text{sum}\rangle$. To do that, we first must introduce our field operators.

State space and basic fields

The field formalism is best explained by example. Suppose we have a single molecular species, B , of interest. We represent a 0-dimensional gas of B monomers as a Fock space with N modes of the B field, $\{B_i\}_{i=1}^N$. Each mode represents a different internal state of the particle. Mathematically each mode acts like a hard-core boson and can be given an explicit representation as a 2D state space (existence or absence of a particle). With a (discrete) internal index ranging from 1 to N , an explicit representation of the vacuum state is

$$|0\rangle = \overbrace{\begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}}^{N \text{ factors}} = \bigotimes_{i=1}^N \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (4.5)$$

Our choice of basis is related to the standard basis for the Pauli matrices. Here the vacuum $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ corresponds to the usual “spin-down.” “Spin-up,” $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, corresponds to the state with a molecule present.

On first impression the internal index resembles the role of a position label on a field in a conventional 1+1-dimensional QFT, but this is not its meaning here. Instead, as its name suggests, the internal index subsumes the internal degrees of freedom which serve to make the complexes classical and (in principle, though not in practice) distinguishable. We will work with a discrete internal index as this avoids some mathematical subtleties. A continuum limit could be taken, but nothing is gained from doing so. For now, our theory is zero-dimensional, in that we are not tracking external position/momenta degrees of freedom of particles. We will speculate later on how position and steric effects might be incorporated.

In the basis we have chosen, operators will be a tensor product similar to Eq. (4.5), but with 2x2 matrices instead of 2x1 vectors. Creation operators for a field B are

defined as

$$\hat{B}_i = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}_i. \quad (4.6)$$

This is shorthand for a tensor product in which the given matrix corresponds to the i -th internal index and all other factors in the tensor product are the identity matrix. When acting on $|0\rangle$, \hat{B}_i leaves the state unchanged except for the vector corresponding to identity index i , which is flipped from $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$. The annihilation operator \check{B}_i is the Hermitian conjugate of \hat{B}_i , and the number operator is $\bar{B}_i = \hat{B}_i \check{B}_i$. We will also find the absense operator, $\tilde{B}_i = \check{B}_i \hat{B}_i$, useful at times. In this basis, \hat{B}_i and \check{B}_i are the usual raising and lowering operators, respectively, in the familiar Pauli basis.

Having given the matrix representation of the basic operators, it is convenient to derive from them all the necessary (anti)commutation relations to avoid ever again needing to write explicit matrix representations. To that end, we tabulate some useful identities and (anti-)commutation relations here for reference:

$$\hat{B}_i \hat{B}_i = \check{B}_i \check{B}_i = 0 \quad (4.7)$$

$$\{\check{B}_i, \hat{B}_i\} = 1 \quad (4.8)$$

$$\bar{B}_i \bar{B}_i = \bar{B}_i \quad (4.9)$$

$$\tilde{B}_i \tilde{B}_i = \tilde{B}_i \quad (4.10)$$

$$[\bar{B}_i, \hat{B}_j] = -[\tilde{B}_i, \hat{B}_j] = \delta_{ij} \hat{B}_j \quad (4.11)$$

$$-[\bar{B}_i, \check{B}_j] = [\tilde{B}_i, \check{B}_j] = \delta_{ij} \check{B}_j \quad (4.12)$$

$$[\check{B}_i, \hat{B}_j] = \delta_{ij} (2\bar{B}_i - 1) = \delta_{ij} (1 - 2\tilde{B}_i). \quad (4.13)$$

As expected, we do not recover the familiar commutation relations from QFT since our particles are neither bosons or fermions. Our relations are, however, similar to some 1D spin chain models, not surprisingly since our Hilbert space is essentially the same.

Internal indices will usually be summed over, so instead of a \hat{B}_i as in Eq. (4.6), it will frequently be convenient to work with the field B defined through

$$\hat{B} = \sum_{i=0}^N \hat{B}_i \quad (4.14)$$

i.e., while \hat{B}_i creates a particle with identity index i , \hat{B} creates a particle with a uniform “superposition” of identities.

From quantum mechanics, one might expect a normalization factor like $1/\sqrt{N}$ to accompany the definition of \hat{B} , but no such factor appears since our particles are classical and distinguishable. We will elaborate on this below. $\hat{B}|0\rangle$ is simply a sum of all possible one particle states. (Note by contrast that $\prod_{i=0}^N \hat{B}_i|0\rangle$ would create a single state containing *many* particles, while Eq. (4.14) creates a superposition of *one* particle states.)

This may become clearer if we consider the sum state for this 0-dimensional gas of monomers, which we are now prepared to do. It is simply the sum of the vacuum, all possible 1-particle states, all possible 2-particle states, etc., given by

$$|\text{sum}\rangle = |0\rangle + \sum_i \hat{B}_i |0\rangle + \frac{1}{2!} \sum_{i,j} \hat{B}_i \hat{B}_j |0\rangle + \cdots = e^{\hat{B}} |0\rangle. \quad (4.15)$$

The factorials correct for overcounting due to permutation symmetry of the sums' dummy indices. But note that we *do* count separately all multi-particle states which differ only in their internal indices. One might be tempted to “correct” this overcounting with a factor of $1/\sqrt{N}$ for each monomer, and although this might appear to work in this trivial case, it would create problems later when we have monomers assembling to form multimers. It is worth noting that only in the $N \rightarrow \infty$ limit do we obtain truly distinguishable particles. Why? Our commutation relations define hard-core bosons, and for finite N we technically have neither Bose-Einstein nor Maxwell-Boltzmann statistics. But if N is sufficiently large compared to the particle numbers relevant in our model, then the mean particle number in each internal mode is low enough that we have classical statistics to a very good approximation. We emphasize that we are not taking the classical limit of a fully-formed quantum theory, so there are no analogies to be made to decoherence theory, Bohr's correspondence principle, or any of the semiclassical approximations commonly used in quantum mechanics.

If we have n non-interacting molecular species of interest, the full state space would be a tensor product of n copies of Eq. (4.5). Fortunately, with the commutation rules laid out, it will not be necessary to work with the explicit matrix representations moving forward.

Binding sites and interactions

Interactions between molecules are neatly handled within the field framework. The simplest example to illustrate is a heterodimer, i.e., two molecular species A and B that bind to each other to form a dimer D . (A homodimer, e.g., binding 2 A

monomers together, is more subtle because of symmetry factors, and is discussed later.) For now, suppose that each molecule has only one binding site, though this is easily generalized below. Suppose also that their internal indices run from 1 to N_A and N_B , respectively (operationally, the internal index serves mainly to distinguish molecules of the same type). We define three new fields to represent this scenario:

- a , which denotes a binding site *on* a molecule of type A that binds *to* a molecule of type B ,
- b (the reverse of a), which denotes a binding site on type B that binds to type A , and
- I , which denotes the interaction itself.

Then a complete dimer is created by the composite operator

$$\hat{D}_{ij} = \hat{A}_i \hat{a}_i \hat{I}_{ij} \hat{b}_j \hat{B}_j, \quad (4.16)$$

and summing over the internal index (which is now a pair (i, j)) gives

$$\hat{D} = \sum_i^{N_A} \sum_j^{N_B} \hat{A}_i \hat{a}_i \hat{I}_{ij} \hat{b}_j \hat{B}_j. \quad (4.17)$$

A nice feature of our formalism is that the internal indices of composite particles naturally emerge from their component parts. This allows one to coarse-grain or fine-grain to whatever level is convenient for the problem at hand and is a significant advantage over existing rule-based methods. In this case, the internal index of a covers the same span as that of A , since they reference the same physical particle, and similarly for b and B . Formally, the vacuum states of a and b are each just another copy of Eq. (4.5), and their creation operators a_i and b_j are identical to those for A_i and B_j , given by Eq. (4.6), but applied to the appropriate Hilbert space. The internal index of I_{ij} runs from 1 to $N_A N_B$, since it must enumerate all possible interaction pairs. Its state space can still be thought of just like Eq. (4.5). The complete state space for this system is simply the tensor product of those for all the various fields: A , B , a , b , and I . Likewise, the internal index of the composite field \hat{D} is now an ordered pair, (i, j) , but this can still be mapped to a 1D state space, Eq. (4.5), just like for monomer fields. Therefore, we can effortlessly coarse-grain the model and regard D itself as an indivisible field, which as we will see later, makes it straightforward to define higher-order complexes such as a tetramer which is itself a dimer of D dimers.

If the site fields appear to be superfluous sidekicks to the interaction fields, it is important to note that they are not. The site fields are necessary to prevent the formation of unphysical complexes, e.g., $\hat{A}_i \hat{B}_j \hat{B}_k \hat{I}_{ij} \hat{I}_{ik}$, representing an A monomer simultaneously bound to two B monomers. Site fields serve to “check” if a binding site is already occupied, since, recalling that each mode acts like a hard-core boson, any attempt to doubly-occupy a mode destroys the whole term. In this case, by pairing the interaction operators with site fields, the analogous expression would be $\hat{A}_i \hat{B}_j \hat{B}_k \hat{a}_i \hat{I}_{ij} \hat{b}_j \hat{a}_i \hat{I}_{ik} \hat{b}_k$, and $\hat{a}_i \hat{a}_i = 0$, removing the entire unphysical complex from consideration. As we will see below, this will have important consequences for designing products of operators to encode desired assembly rules, as poorly designed rules can lead to unphysical complexes of this sort.

“Wick-like” results

Having surveyed the behavior of the state space and basic fields of our formalism, it will be useful next to compute some standard results that will appear repeatedly in the calculations that follow.

Wick’s Theorem does not apply since we have neither bosons nor fermions, but some very similar and equally useful results can be seen directly from the commutation relations. In evaluating the partition function we will encounter expressions containing strings annihilation operators followed by creation operators, like $\langle 0 | \check{B}_i \check{B}_j \hat{B}_k \hat{B}_l | 0 \rangle$. As with Wick’s theorem, since annihilation operators annihilate $|0\rangle$, the goal is to permute them to the right of an expression. The simplest prototype is $\check{B}_i \hat{B}_j | 0 \rangle$. To resolve this, use the commutation relations and above identities to find

$$\langle 0 | \check{B}_i \hat{B}_j | 0 \rangle = \langle 0 | (\hat{B}_j \check{B}_i + \delta_{ij} (1 - 2\bar{B}_j)) | 0 \rangle = \delta_{ij}. \quad (4.18)$$

In this case, the annihilation and number operators were already adjacent to the vacuum and could be resolved. In more complicated expressions, there may be additional creation operators in the way. But \bar{B}_j can always be commuted through them as its internal index must be different (otherwise the whole expression would vanish since $\hat{B}_j \hat{B}_j = 0$).

We will refer to Eq. (4.18) as a *contraction*; though it does not quite match the standard definition, we wish to elicit an association as the idea and usage are virtually identical.

Longer strings are evaluated by repeatedly applying the commutation relations. With Eq. (4.18) we can evaluate longer strings by simply contracting all possible

combinatorial pairings like

$$\langle 0 | \check{B}_i \check{B}_j \hat{B}_k \hat{B}_l | 0 \rangle = \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk} \quad (4.19)$$

or

$$\begin{aligned} \langle 0 | \check{B}_i \check{B}_j \check{B}_k \hat{B}_l \hat{B}_m \hat{B}_n | 0 \rangle &= \delta_{il} (\delta_{jm} \delta_{kn} + \delta_{jn} \delta_{km}) \\ &\quad + \delta_{im} (\delta_{jl} \delta_{kn} + \delta_{jn} \delta_{kl}) \\ &\quad + \delta_{in} (\delta_{jl} \delta_{km} + \delta_{jm} \delta_{kl}) \end{aligned} \quad (4.20)$$

Clearly, if there are q creation operators following an equal number of annihilation operators, there will be $q!$ terms (the number of permutations on q objects). If the number of operators is not equal, the entire term will vanish (the operators left over after all possible contractions will annihilate one vacuum or the other).

There are some subtleties with multiply-indexed fields, e.g., interaction fields. We can think of interactions as *directed* or *undirected*. The intuitive distinction is, roughly, whether the two binding targets of the interaction are distinguishable. However, the distinction could for instance be controlled by the interaction field's attendant binding sites, rather than the binding targets themselves being distinguishable. If the interaction is directed, the indices are *not* interchangeable and the respective indices must contract, i.e.,

$$\langle 0 | \check{I}_{ij} \hat{I}_{kl} | 0 \rangle = \delta_{ik} \delta_{jl}, \quad (4.21)$$

$$\langle 0 | \check{I}_{ij} \check{I}_{kl} \hat{I}_{mn} \hat{I}_{pq} | 0 \rangle = \delta_{im} \delta_{jn} \delta_{kp} \delta_{lq} + \delta_{ip} \delta_{jq} \delta_{km} \delta_{ln}, \quad (4.22)$$

and similarly if the interaction field carries more than two indices. If, however, the interaction is an undirected type, we must consider all equivalent permutations, such as

$$\langle 0 | \check{\lambda}_{ij} \check{\lambda}_{kl} | 0 \rangle = \delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk} \quad (4.23)$$

$$\begin{aligned} \langle 0 | \check{\lambda}_{ijk} \hat{\lambda}_{lmn} | 0 \rangle &= \delta_{il} (\delta_{jm} \delta_{kn} + \delta_{jn} \delta_{km}) \\ &\quad + \delta_{im} (\delta_{jl} \delta_{kn} + \delta_{jn} \delta_{kl}) + \delta_{in} (\delta_{jl} \delta_{km} + \delta_{jm} \delta_{kl}) \end{aligned} \quad (4.24)$$

$$\begin{aligned} \langle 0 | \check{\lambda}_{ij} \check{\lambda}_{kl} \hat{\lambda}_{mn} \hat{\lambda}_{pq} | 0 \rangle &= (\delta_{im} \delta_{jn} + \delta_{in} \delta_{jm}) (\delta_{kp} \delta_{lq} + \delta_{kq} \delta_{lp}) \\ &\quad + (\delta_{ip} \delta_{jq} + \delta_{iq} \delta_{jp}) (\delta_{km} \delta_{ln} + \delta_{kn} \delta_{lm}). \end{aligned} \quad (4.25)$$

There are several points to note here:

- Although it may seem obvious, Eq. (4.23) hides some subtleties. Essentially, we are defining $\lambda_{ij} \equiv \lambda_{ji}$ for indistinguishable interactions. It might sound

preferable to simply not allow $j < i$ which would avoid overcounting the state space. But this would require awkwardly restricting sums over internal indices and be extremely inconvenient in practice.

- Note the similarity between Eq. (4.23) and Eq. (4.19), and between Eq. (4.24) and Eq. (4.20). Single multi-indexed operators with interchangeable indices behave simply like a product of single indexed fields. This is not as useful as it sounds because we will rarely encounter expressions like Eq. (4.23) without also encountering ones like Eq. (4.25) arising from the interaction's attendant site fields.
- The new behavior occurs when we have more than one creation and annihilation operator. Reasoning from Eq. (4.23) and Eq. (4.24), one might expect the right hand side of Eq. (4.25) to behave like four annihilation and four creation operators, all single indexed. But this is incorrect: it neglects the physical fact that the interactions cannot be “split;” if i is paired with m , j cannot be paired with p or q . As a result, there are only 8, not 24 terms produced.
- One could obviously construct mixed cases as well, where some of the interaction's indices are interchangeable and some are not.
- Since these formulae present a *much* more stringent constraint on allowed internal indices than do the single-index fields that (nearly always) accompany these interaction fields, it is usually convenient to analyze and contract the interaction fields first.

Now that we have a library of basic results, we move on to consider the internal index-summed fields we will usually be using, i.e., summed over internal indices as in Eq. (4.14). In this case the above results retain a similar form. Formally, if $q < N_A$,

$$\begin{aligned}
 \langle 0 | (\check{A})^p (\hat{A})^q | 0 \rangle &= \langle 0 | \sum_{i_1, \dots, i_p} \sum_{j_1, \dots, j_q} \left(\prod_{\alpha=1}^p \check{A}_{i_\alpha} \right) \left(\prod_{\beta=1}^q \hat{A}_{j_\beta} \right) | 0 \rangle \\
 &= \delta_{pq} q! \frac{N_A!}{(N_A - q)!} \approx \delta_{pq} q! (N_A)^q.
 \end{aligned} \tag{4.26}$$

Intuitively, why is this result true? As above, the expression vanishes if $p \neq q$. If $p = q$, then there are $N_A(N_A - 1) \cdots (N_A - q + 1)$ possible choices for the internal indices, giving the factor of $\frac{N_A!}{(N_A - q)!}$. Note that if $q < N_A$, the expression vanishes

because at least one internal index is duplicated, and $\hat{A}_i \hat{A}_i = 0$. But insisting on $q > N_A$ is not enough. Only in the limit of $N_A \gg q$ does this factor reduce to $(N_A)^q$, which illustrates why taking the number of internal indices to infinity is necessary.¹ For any particular choice of internal indices, we carry out the q contractions exactly as above, producing $q!$ terms arising from permutations.

We will also frequently encounter expressions involving a product of number operators, followed by creation operators, acting on $|0\rangle$, i.e.,

$$(\bar{A})^p (\hat{A})^q |0\rangle = \sum_{i_1, \dots, i_p}^{N_A} \sum_{j_1, \dots, j_q}^{N_A} \left(\prod_{\alpha=1}^p \bar{A}_{i_\alpha} \right) \left(\prod_{\beta=1}^q \hat{A}_{j_\beta} \right) |0\rangle \quad (4.27)$$

To evaluate this, note that $\hat{A}_i |0\rangle$ is an eigenstate of \bar{A}_i (with eigenvalue 1). Consider a single term in the sum. There are q choices for the internal index of each number operator that will not cause the term to vanish, and these choices are independent (unlike in Eq. (4.26), multiple number operators may “contract” with the same creation operator). Put another way, i_1 could equal any of j_1, j_2, \dots, j_q , and since A_{i_1} is acting on an eigenstate, we have the same choices for i_2 , and so on. Therefore,

$$\begin{aligned} (\bar{A})^p (\hat{A})^q |0\rangle &= \sum_{i_1, \dots, i_p}^{N_A} \sum_{j_1, \dots, j_q}^{N_A} \left(\prod_{\alpha=1}^p \delta_{i_\alpha j_\beta} \right) \left(\prod_{\beta=1}^q \hat{A}_{j_\beta} \right) |0\rangle \\ &= q^p \sum_{j_1, \dots, j_q}^{N_A} \left(\prod_{\beta=1}^q \hat{A}_{j_\beta} \right) |0\rangle \\ &= q^p (\hat{A})^q |0\rangle. \end{aligned} \quad (4.28)$$

Note as for Eq. (4.26) above, this result is only true if $q < N_A$, and vanishes otherwise. Intuitively, each of the p number operators could “detect” any one of the q fields, so there are q^p ways of arranging the pairing. This result is useful in evaluating expressions such as $\exp(cA) \exp(\hat{A}) |0\rangle$ occurring, among other places, when we evaluate partition functions, where c is an arbitrary real number. This

¹ Without this simplification, Eq. (4.30) below would not simplify. Our trick described later in the Hamiltonian would no longer cancel away the dependence on internal indices, and this dependence on internal indices would propagate through to the partition function, which is obviously undesirable.

works out to

$$\begin{aligned}
 \exp(cA) \exp(\hat{A}) |0\rangle &= \sum_{pq} \frac{(c\bar{A})^p}{p!} \frac{(\hat{A})^q}{q!} |0\rangle \\
 &= \sum_{pq} \frac{(cq)^p}{p!} \frac{(\hat{A})^q}{q!} |0\rangle \\
 &= \sum_q (e^c)^q \frac{(\hat{A})^q}{q!} |0\rangle \\
 &= \exp(e^c \hat{A}) |0\rangle.
 \end{aligned} \tag{4.29}$$

If we then want the inner product with $\langle 0 | \exp(\check{A})$, we find, using Eq. (4.26)

$$\begin{aligned}
 \langle 0 | \exp(\check{A}) \exp(e^c \hat{A}) |0\rangle &= \sum_{pq} e^{cq} \langle 0 | \frac{(\check{A})^p}{p!} \frac{(\hat{A})^q}{q!} |0\rangle \\
 &= \sum_{pq} e^{cq} \frac{\delta_{pq} q!}{p! q!} \frac{N_A!}{(N_A - q)!} \\
 &\approx \sum_q e^{cq} \frac{(N_A)^q}{q!} \\
 &= \exp(e^c N_A),
 \end{aligned} \tag{4.30}$$

again taking the limit $q \ll N_A$ in the second to last line.

The Gallery

With the basics of the formalism in hand, we now turn to the construction of the $|\text{sum}\rangle$ state. We will sketch the general result and demonstrate particular examples in subsequent sections.

Suppose we have a complete list of all possible complexes that may be formed in a particular system, and denote the creation operators for these complexes as $\hat{A}, \hat{B}, \hat{C}, \dots$. Define \mathbb{G} as the sum of these operators, i.e.,

$$\mathbb{G} = \hat{A} + \hat{B} + \hat{C} + \dots \tag{4.31}$$

Note that some of these operators may themselves be composite particles, but even these contain only creation operators and no number, absence or annihilation operators, like the dimer in 4.17. Therefore they all commute with each other.

If we can construct \mathbb{G} , we can use it to construct the $|\text{sum}\rangle$ state as

$$\begin{aligned} |\text{sum}\rangle &= \exp(\mathbb{G}) |0\rangle = \exp(\hat{A} + \hat{B} + \hat{C} + \dots) |0\rangle \\ &= \exp(\hat{A}) \exp(\hat{B}) \exp(\hat{C}) \dots |0\rangle \\ &= \left(\sum_p \frac{(\hat{A})^p}{p!} \right) \left(\sum_q \frac{(\hat{B})^q}{q!} \right) \left(\sum_r \frac{(\hat{C})^r}{r!} \right) \dots |0\rangle. \end{aligned} \quad (4.32)$$

The n -th term in each series produces all possible states with n copies of that molecule. The factorial corrects for overcounting from internal indices. For instance,

$$\frac{(\hat{A})^p}{p!} = \frac{1}{p!} \sum_{k_1, \dots, k_p} \hat{A}_{k_1} \hat{A}_{k_2} \dots \hat{A}_{k_p} = \sum_{k_1 < k_2 < \dots < k_p} \hat{A}_{k_1} \hat{A}_{k_2} \dots \hat{A}_{k_p}, \quad (4.33)$$

where the k_i are the identity indices for the various copies of the field. The unconstrained multiple sum produces p copies of type A molecules, but overcounts them $p!$ ways. After dividing by $p!$, the entire term produces each distinguishable state with p particles exactly once. Then multiplying the series for A, B, C, \dots , clearly produces states with all possible copy numbers of all complexes. The reader may be alarmed that this state is not normalized in any reasonable sense, i.e., $\langle \text{sum} | 1 | \text{sum} \rangle$ is not finite. But the notational analogy with familiar quantum mechanics is misleading here: $|\text{sum}\rangle$ is not itself a physical state, but rather it catalogs all possible pure physical states by summing up their operator representations. Just as in standard formulations of equilibrium statistical mechanics, we do not require the enumeration of microstates to be finite or “normalizable” in any sense. Only after weighting states by their corresponding Boltzmann weights should we hope for a finite number, and the same is true here.

From the above sketch, it should be plausible that $|\text{sum}\rangle = e^{\mathbb{G}} |0\rangle$ holds for any system if \mathbb{G} is the sum of creation operators of all possible complexes, summed over all internal states. Doi made a similar observation [9], but the operator \mathbb{G} , does not yet seem to have a name. We will call \mathbb{G} the “gallery,” evoking a buffet of all possible complexes available to the system.

This tells how to construct $|\text{sum}\rangle$, but this procedure is unwieldy if the system permits more than a few complexes. It still requires us to enumerate every possible complex in the system, which is exactly what we want to avoid.

The Hamiltonian

Now that we have seen the (inefficient) construction of the $|\text{sum}\rangle$ state, we turn to specifying the Hamiltonian, with which we will have all we need to calculate the partition function. In fact, the construction of the Hamiltonian will help motivate the efficient construction of $|\text{sum}\rangle$ in the next section.

We constructed the gallery as a sum of all possible complexes, composed out of creation operators. The qualitative idea of the Hamiltonian is to “detect” features of those complexes using number operators and assign a free energy to each feature. Then the total free energy of a complex is the sum of the free energies of all its features. Features of interest might include everything from the chemical potential of a monomer, pairwise binding interactions, the allosteric state of a molecule, on up to a complicated cooperativity energy contingent on a certain set of monomers being present in a complex. An example may clarify.

Consider again the heterodimer example discussed earlier. Allowed complexes in this model are monomers A and B along with an AB dimer. An explicit gallery for this model is

$$\mathbb{G} = \hat{A} + \hat{B} + \hat{D} = \sum_i^{N_A} \hat{A}_i + \sum_j^{N_B} \hat{B}_j + \sum_k^{N_A} \sum_l^{N_B} \hat{A}_k \hat{a}_k \hat{I}_{kl} \hat{b}_l \hat{B}_l. \quad (4.34)$$

A reasonable Hamiltonian for this model would include chemical potentials μ_A and μ_B for the monomer species and a binding free energy ϵ for the dimer interaction. Encoded in our operators, this has the form

$$\mathbb{H} = \mu'_A \sum_i^{N_A} \bar{A}_i + \mu'_B \sum_j^{N_B} \bar{B}_j + \epsilon \sum_k^{N_A} \sum_l^{N_B} \bar{a}_k \bar{I}_{kl} \bar{b}_l. \quad (4.35)$$

where the “renormalized” chemical potentials are $\mu'_A = \mu_A - k_B T \ln N_A$ and $\mu'_B = \mu_B - k_B T \ln N_B$, and energy parameters are nondimensionalized by $k_B T$ for convenience. Note that the unprimed μ s are the real, physical chemical potentials. The reader may rightly wonder why this renormalization is necessary. Recall that in the construction of the $|\text{sum}\rangle$ state, we counted separately states which differed only in their internal indices. In the present example, for instance, there is a single state with zero particles, the vacuum $|0\rangle$, there are N_A states with a single A monomer, $N_A \times N_B$ states with one A and one B , $N_A \times N_B^2$ states with one B and one AB , etc. We do not actually wish to count these states as distinguishable, and renormalizing the chemical potentials of monomers is the simplest method of “fixing” this.

Redefining field operators as, e.g., $\hat{A} = \frac{1}{\sqrt{N_A}} \sum_i^{N_A} \hat{A}_i$, may sound simpler and appealing by analogy to quantum mechanics. We explored this idea and found it leads to confusingly inconsistent rules, where certain monomer fields in the gallery require factors of $1/\sqrt{N}$, certain composite fields do not, analogous operators in the Hamiltonian do not, and exceptions and special cases seemed to mount with every model we explored. It also spoiled the natural coarse-graining of composite particles, as counting factors did not always match up with their internal indices, which is a feature of our method we would like to retain.

In fact, renormalizing the chemical potentials is not so opaque or arbitrary as it may seem. An internal index running from 1 to N in a sense adds N orthogonal microstates to the system: useful because they render our particles effectively classical rather than quantum objects, but inconvenient for our enumeration of states. So $k_B T \ln N$ looks like T times the extra entropy associated with the “fictitious” internal index degrees of freedom that we introduced. This also explains why our attempts to introduce corrective factors in the definition of index-averaged fields failed: the miscounting is not associated with particular fields but rather with the number of unique indices in a complex. As we will see later, for more complex models when several operators may act on a complex in sequence, it would become quite tedious to work out in advance which indices will or will not interact and therefore where $1/\sqrt{N}$ factors are or are not needed.

Finally, note that our operator for detecting the binding interaction associated with energy ϵ detected both the interaction field I as well as the site fields a and b . In this simple example this is redundant and not strictly necessary, but as models grow in complexity we have found it generally advisable to construct Hamiltonian operators as specifically as possible to ensure they trigger only on the desired complexes and no others unintentionally. For instance, when several species are present and each has multiple binding partners, potentially with multiple binding sites on some or all of those partners, it is easy to overlook and fail to exclude an unintended binding target.

The factory

The construction of the Hamiltonian with our operator formalism suggests a way to avoid having to construct the complete gallery in order to construct $s|\text{sum}\rangle$. If we can design operators in the Hamiltonian that detect specific pieces of a complex, why not do the same for construction? In other words, can we detect existing

features but then, rather than merely assigning a free energy, graft on new features, thereby assembling a complex piece-by-piece according to a set of rules rather than constructing it at a stroke as the gallery requires?

This is in fact possible. We dub this object the *factory*, churning out complexes like an assembly line in contrast to the gallery, a collection of fully-formed objects. For a given model, the factory is defined as an ordered list of operators, such as

$$\mathcal{F} = (\mathbb{F}_1, \mathbb{F}_2, \dots, \mathbb{F}_n), \quad (4.36)$$

which produces the sum state according to

$$|\text{sum}\rangle = e^{\mathbb{F}_n} \dots e^{\mathbb{F}_2} e^{\mathbb{F}_1} |0\rangle. \quad (4.37)$$

The order of the \mathbb{F}_i is essential. It is also essential that the contents of each \mathbb{F}_i commute, but they need not commute with each other, e.g., \mathbb{F}_1 can be a sum of products of operators that all commute, and separately all the operators in \mathbb{F}_2 must commute, but $[\mathbb{F}_1, \mathbb{F}_2]$ need not equal 0. In fact the lack of commutation between the various \mathbb{F}_i in \mathcal{F} is what allows the factory to generate nontrivial complexes. Requiring the contents of each \mathbb{F}_i to commute ensures there is no ambiguity in expanding $e^{\mathbb{F}_i}$.

To emphasize this another way, note that we did *not* define the factory as a sum analogous to the gallery, e.g., a definition such as

$$|\text{sum}\rangle \neq e^{\mathbb{F}_n + \dots + \mathbb{F}_2 + \mathbb{F}_1} |0\rangle. \quad (4.38)$$

does not work since the exponential cannot be expanded unambiguously (unless all the \mathbb{F}_i commute, in which case we have simply reproduced the gallery). One might imagine defining rules or conventions for how to resolve the ordering of operators in Eq. (4.38), but such procedures inevitably become more far more complicated than the recipe of Eq. (4.36) and Eq. (4.37).

One critical consideration in designing a factory for a given model is that the factory must generate every finished complex in the model in exactly one and only one way, via a single assembly path. Lest the reader think that this error merely produces a different, yet still valid, model from the one intended, this is not the case. If different complexes can be generated in different numbers of ways, this spoils the equal statistical weight of all the complexes in $|\text{sum}\rangle$, and it follows that $|\text{sum}\rangle \neq e^{\mathbb{F}_n} \dots e^{\mathbb{F}_1} |0\rangle$. So this is not merely a different model, it is a nonsensical model. We will see examples of how this may occur below.

References

- ¹M. Lynch, “The evolution of multimeric protein assemblages”, *Molecular Biology and Evolution* **29**, 1353–1366 (2012).
- ²J. Monod, J. Wyman, and J. P. Changeux, “On the nature of allosteric transitions: A plausible model”, *Journal of Molecular Biology* **12**, 88–118 (1965).
- ³J.-P. Changeux, “Allostery and the monod-wyman-changeux model after 50 years”, *Annual Review of Biophysics* **41**, 103–133 (2012).
- ⁴V. N. Kampen, “Stochastic processes in physics and chemistry”, in *Stoch. Process. Phys. Chem.* 3rd (2007).
- ⁵A. Sanchez and J. Kondev, “Transcriptional control of noise in gene expression”, *Proceedings of The National Academy Of Sciences Of The United States Of America* **105**, 5081–6 (2008).
- ⁶A. Sanchez, H. G. Garcia, D. Jones, R. Phillips, and J. Kondev, “Effect of promoter architecture on the cell-to-cell variability in gene expression”, *PLoS computational biology* **7**, e1001100 (2011).
- ⁷E. L. King and C. Altman, “A Schematic Method of Deriving the Rate Laws for Enzyme-Catalyzed Reactions”, *The Journal of Physical Chemistry* **60**, 1375–1378 (1956).
- ⁸T. L. Hill, “Studies in irreversible thermodynamics IV. diagrammatic representation of steady state fluxes for unimolecular systems”, *Journal of Theoretical Biology* **10**, 442–459 (1966).
- ⁹M. Doi, “Second quantization representation for classical many-particle system”, *J. Phys. A. Math. Gen.* **9**, 1465–1477 (1976).
- ¹⁰M. Doi, “Stochastic theory of diffusion-controlled reaction”, *J. Phys. A. Math. Gen.* **9**, 1479–1495 (1976).
- ¹¹P. Grassberger and M. Scheunert, “Fock space methods for identical classical objects”, *Fortschritte der Phys.* **28**, 547–578 (1980).
- ¹²S. C. Park and J. M. Park, “Generating function, path integral representation, and equivalence for stochastic exclusive particle systems”, *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **71**, 026113 (2005).
- ¹³L. Peliti, “Path integral approach to birth-death processes on a lattice”, *Journal de Physique* **46**, 1469–1483 (1985).
- ¹⁴A. M. Walczak, A. Mugler, and C. H. Wiggins, “Analytic methods for modeling stochastic regulatory networks”, *arXiv*, 1005.2648v1 (2010).
- ¹⁵D. C. Mattis and M. L. Glasser, “The uses of quantum field theory in diffusion-limited reactions”, *Reviews of Modern Physics* **70**, 979–1001 (1998).
- ¹⁶P. J. Dodd and N. M. Ferguson, “A many-body field theory approach to stochastic models in population biology”, *PLoS One* **4**, e6855 (2009).

- ¹⁷U. C. Täuber, “Population oscillations in spatial stochastic Lotka–Volterra models: a field-theoretic perturbational analysis”, *J. Phys. A Math. Theor.* **45**, 405002 (2012).
- ¹⁸B. Zhang and P. G. Wolynes, “Stem cell differentiation as a many-body problem”, *Proceedings of the National Academy of Sciences* **111**, 10185–10190 (2014).
- ¹⁹M. A. Buice and J. D. Cowan, “Field-theoretic approach to fluctuation effects in neural networks”, *Physical Review E* **75**, 051919 (2007).
- ²⁰J. D. Cowan, J. Neuman, B. Kiewiet, and W. van Drongelen, “Self-organized criticality in a network of interacting neurons”, *Journal of Statistical Mechanics: Theory and Experiment* **2013**, P04030 (2013).
- ²¹P. C. Bressloff and J. M. Newby, “Path integrals and large deviations in stochastic hybrid systems”, *Physical Review E* **89**, 042701 (2014).
- ²²L. Cui, I. Murchland, K. E. Shearwin, and I. B. Dodd, “Enhancer-like long range transcriptional activation by λ CI-mediated DNA looping”, *Proceedings of the National Academy of Sciences* **110**, 2922–2927 (2013).
- ²³M. Razo-Mejia, J. Q. Boedicker, D. Jones, A. DeLuna, J. B. Kinney, and R. Phillips, “Comparison of the theoretical and real-world evolutionary potential of a genetic circuit”, *Physical Biology* **11**, 026005 (2014).
- ²⁴W. S. Hlavacek, J. R. Faeder, M. L. Blinov, A. S. Perelson, and B. Goldstein, “The complexity of complexes in signal transduction”, *Biotechnology and Bioengineering* **84**, 783–94 (2003).
- ²⁵W. S. Hlavacek, J. R. Faeder, M. L. Blinov, R. G. Posner, M. Hucka, and W. Fontana, “Rules for modeling signal-transduction systems”, *Science’s STKE*, **10**. 1126/stke.3442006re6 (2006).
- ²⁶L. A. Chylek, L. A. Harris, J. R. Faeder, and W. S. Hlavacek, “Modeling for (physical) biologists: An introduction to the rule-based approach”, *Physical Biology* **12**, 045007 (2015).
- ²⁷G. Berry and G. Boudol, “The chemical abstract machine”, *Theoretical Computer Science* **96**, 217–248 (1992).
- ²⁸A. Regev, E. M. Panina, W. Silverman, L. Cardelli, and E. Shapiro, “BioAmbients: An abstraction for biological compartments”, *Theoretical Computer Science* **325**, 141–167 (2004).
- ²⁹V. Danos and C. Laneve, “Formal molecular biology”, *Theoretical Computer Science* **325**, 69–110 (2004).
- ³⁰J. R. Faeder, M. L. Blinov, B. Goldstein, and W. S. Hlavacek, “Rule-based modeling of biochemical networks”, *Complexity* **10**, 22–41 (2005).
- ³¹G. Boniolo, M. D’Agostino, and P. P. Di Fiore, “Zsyntax: A formal language for molecular biology with projected applications in text mining and biological prediction”, *PLoS One* **5**, e9511 (2010).

- ³²J. Yang, M. I. Monine, J. R. Faeder, and W. S. Hlavacek, “Kinetic Monte Carlo method for rule-based modeling of biochemical networks”, *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* **78**, 031910 (2008).
- ³³J. R. Faeder, M. L. Blinov, and W. S. Hlavacek, “Rule-based modeling of biochemical systems with BioNetGen”, *Methods in Molecular Biology* **500**, 113–67 (2009).
- ³⁴J. Feret, V. Danos, J. Krivine, R. Harmer, and W. Fontana, “Internal coarse-graining of molecular systems”, *Proceedings of the National Academy of Sciences of the United States of America* **106**, 6453–6458 (2009).
- ³⁵M. W. Sneddon and T. Emonet, “Modeling cellular signaling: Taking space into the computation”, *Nature Methods* **9**, 239–42 (2012).
- ³⁶A. M. Smith, W. Xu, Y. Sun, J. R. Faeder, and G. E. Marai, “RuleBender: integrated modeling, simulation and visualization for rule-based intracellular biochemistry”, *BMC Bioinformatics* **13**, S3 (2012).

Appendix A

APPENDICES FOR CHAPTER II, RECONCILING KINETIC AND EQUILIBRIUM MODELS OF BACTERIAL TRANSCRIPTION

A.1 Derivations for non-bursty promoter models

In this section we detail the calculation of mean mRNA levels, fold-changes in expression, and Fano factors for nonequilibrium promoter models 1 through 4 in Figure 2.1. These are the results that were quoted but not derived in Sections 2.3 and 2.4 of the main text. In each of these four models, the natural mathematicization of their cartoons is as a chemical master equation such as Eq. 2.12 for model 1. Before jumping into the derivations of the general computation of the mean mRNA level and the Fano factor we will work through the derivation of an example master equation. In particular we will focus on model 1 from Figure 2.1(C). The general steps are applicable to all other chemical master equations in this work.

Derivation of chemical master equation




Figure A.1: **Two-state promoter chemical master equation.** (A) Schematic of the two state promoter simple repression model. Rates k_R^+ and k_R^- are the association and dissociation rates of the transcriptional repressor, respectively, r is the transcription initiation rate, and γ is the mRNA degradation rate. (B) Schematic depiction of the mRNA count state transitions. The model in (A) only allows for jumps in mRNA of size 1. The production of mRNA can only occur when the promoter is in the transcriptionally active state.

The chemical master equation describes the continuous time evolution of a continuous or discrete probability distribution function. In our specific case we want to describe the time evolution of the discrete mRNA distribution. What this means is that we want to compute the probability of having m mRNA molecules at time $t + \Delta t$, where Δt is a sufficiently small time interval such that only one of the possible reactions take place during that time interval. For the example that we will work out here in detail we chose the two-state stochastic simple repression model schematized in Figure A.1(A). To derive the master equation we will focus more on the representation shown in Figure A.1(B), where the transitions between different mRNA counts and promoter states is more explicitly depicted. Given that the DNA promoter can exist in one of two states – transcriptionally active state, and with repressor bound – we will keep track not only of the mRNA count, but on which state the promoter is. For this we will keep track of two probability distributions: The probability of having m mRNAs at time t when the promoter is in the transcriptionally active state A , $p_A(m, t)$, and the equivalent probability but when the promoter is in the repressor bound state R , $p_R(m, t)$.

Since mRNA production can only take place in the transcriptionally active state we will focus on this function for our derivation. The repressor bound state will have an equivalent equation without terms involving the production of mRNAs. We begin by listing the possible state transitions that can occur for a particular mRNA count with the promoter in the active state. For state changes in a small time window Δt that “jump into” state m in the transcriptionally active state we have

- Produce an mRNA, jumping from $m - 1$ to m .
- Degrade an mRNA, jumping from $m + 1$ to m .
- Transition from the repressor bound state R with m mRNAs to the active state A with m mRNAs.

Likewise, for state transitions that “jump out” of state m in the transcriptionally inactive state we have

- Produce an mRNA, jumping from m to $m + 1$.
- Degrade an mRNA, jumping from m to $m - 1$.
- Transition from the active state A with m mRNAs to the repressor bound state R with m mRNAs.

The mRNA production does not depend on the current number of mRNAs, therefore these state transitions occur with probability $r\Delta t$. The same is true for the promoter state transitions; each occurs with probability $k_R^\pm\Delta t$. As for the mRNA degradation events, these transitions depend on the current number of mRNA molecules since the more molecules of mRNA there are, the more will decay during a given time interval. Each molecule has a constant probability $\gamma\Delta t$ of being degraded, so the total probability for an mRNA degradation event to occur is computed by multiplying this probability by the current number of mRNAs.

To see these terms in action let us compute the probability of having m mRNA at time $t + \Delta t$ in the transcriptionally active state. This takes the form

$$\begin{aligned}
 p_A(m, t + \Delta t) = & p_A(m, t) \\
 & + \overbrace{(r\Delta t)p_A(m-1, t)}^{m-1 \rightarrow m} - \overbrace{(r\Delta t)p_A(m, t)}^{m \rightarrow m+1} \\
 & + \overbrace{(m+1)(\gamma\Delta t)p_A(m+1, t)}^{m+1 \rightarrow m} - \overbrace{m(\gamma\Delta t)p_A(m, t)}^{m \rightarrow m-1} \\
 & + \overbrace{(k_R^-\Delta t)p_R(m, t)}^{R \rightarrow A} - \overbrace{(k_R^+\Delta t)p_A(m, t)}^{A \rightarrow R},
 \end{aligned} \tag{A.1}$$

where the overbrace indicates the corresponding state transitions. Notice that the second to last term on the right-hand side is multiplied by $p_R(m, t)$ since the transition from state R to state A depends on the probability of being in state R to begin with. It is through this term that the dynamics of the two probability distribution functions ($p_R(m, t)$ and $p_A(m, t)$) are coupled. An equivalent equation can be written for the probability of having m mRNA at time $t + \Delta t$ while in the repressor bound state, the only difference being that the mRNA production rates are removed, and the sign for the promoter state transitions are inverted. This is

$$\begin{aligned}
 p_R(m, t + \Delta t) = & p_R(m, t) \\
 & + \overbrace{(m+1)(\gamma\Delta t)p_R(m+1, t)}^{m+1 \rightarrow m} - \overbrace{m(\gamma\Delta t)p_R(m, t)}^{m \rightarrow m-1} \\
 & - \overbrace{(k_R^-\Delta t)p_R(m, t)}^{R \rightarrow A} + \overbrace{(k_R^+\Delta t)p_A(m, t)}^{A \rightarrow R}.
 \end{aligned} \tag{A.2}$$

All we have to do now are simple algebraic steps in order to simplify the equations. Let us focus on the transcriptionally active state A . First we will send the term

$p_A(m, t)$ to the right-hand side, and then we will divide both sides of the equation by Δt . This results in

$$\begin{aligned} \frac{p_A(m, t + \Delta t) - p_A(m, t)}{\Delta t} = & r p_A(m - 1, t) - r p_A(m, t) \\ & + (m + 1) \gamma p_A(m + 1, t) - m \gamma p_A(m, t) \\ & + k_R^- p_R(m, t) - k_R^+ p_A(m, t). \end{aligned} \quad (\text{A.3})$$

Upon taking the limit when $\Delta t \rightarrow 0$ we can transform the left-hand side into a derivative, obtaining the chemical master equation

$$\begin{aligned} \frac{dp_A(m, t)}{dt} = & r p_A(m - 1, t) - r p_A(m, t) \\ & + (m + 1) \gamma p_A(m + 1, t) - m \gamma p_A(m, t) \\ & + k_R^- p_R(m, t) - k_R^+ p_A(m, t). \end{aligned} \quad (\text{A.4})$$

Doing equivalent manipulations for the repressor bound state gives an ODE of the form

$$\begin{aligned} \frac{dp_R(m, t)}{dt} = & (m + 1) \gamma p_R(m + 1, t) - m \gamma p_R(m, t) \\ & - k_R^- p_R(m, t) + k_R^+ p_A(m, t). \end{aligned} \quad (\text{A.5})$$

In the next section we will write these coupled ODEs in a more compact form using matrix notation.

Matrix form of the multi-state chemical master equation

Having derived an example chemical master equation we now focus on writing a general matrix form for the kinetic models 1-4 shown in Figure 2.1(C) in the main text. In each of these four models, the natural mathematicization of their cartoons is as a chemical master equation such as Eq. 2.12 for model 1, which for completeness we reproduce here as

$$\begin{aligned} \frac{d}{dt} p_R(m, t) = & \overbrace{-k_R^- p_R(m, t)}^{R \rightarrow U} + \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{(m + 1) \gamma p_R(m + 1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_R(m, t)}^{m \rightarrow m-1} \\ \frac{d}{dt} p_U(m, t) = & \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} - \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{r p_U(m - 1, t)}^{m-1 \rightarrow m} - \overbrace{r p_U(m, t)}^{m \rightarrow m+1} \\ & + \overbrace{(m + 1) \gamma p_U(m + 1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_U(m, t)}^{m \rightarrow m-1}. \end{aligned} \quad (\text{A.6})$$

Here $p_R(m, t)$ and $p_U(m, t)$ are the probabilities of finding the system with m mRNA molecules at time t either in the repressor bound or unbound states, respectively. r

is the probability per unit time that a transcript will be initiated when the repressor is unbound, and γ is the probability per unit time for a given mRNA to be degraded. k_R^- is the probability per unit time that a bound repressor will unbind, while k_R^+ is the probability per unit time that an unbound operator will become bound by a repressor. Assuming mass action kinetics, k_R^+ is proportional to repressor copy number R .

Next consider the cartoon for nonequilibrium model 2 in Figure 2.1(C). Now we must track probabilities p_R , p_P , and p_E for the repressor bound, empty, and polymerase bound states, respectively. By analogy to Eq. A.6, the master equation for model 2 can be written

$$\begin{aligned}
 \frac{d}{dt}p_R(m, t) &= - \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} + \overbrace{k_R^+ p_E(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_R(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_R(m, t)}^{m \rightarrow m-1} \\
 \frac{d}{dt}p_E(m, t) &= \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} - \overbrace{k_R^+ p_E(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_E(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_E(m, t)}^{m \rightarrow m-1} . \\
 &\quad + \overbrace{k_P^- p_P(m, t)}^{A \rightarrow U} - \overbrace{k_P^+ p_E(m, t)}^{U \rightarrow A} + \overbrace{r p_P(m-1, t)}^{m-1 \rightarrow m, A \rightarrow U} \\
 \frac{d}{dt}p_P(m, t) &= - \overbrace{k_P^- p_P(m, t)}^{A \rightarrow U} + \overbrace{k_P^+ p_E(m, t)}^{U \rightarrow A} + \overbrace{(m+1)\gamma p_P(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_P(m, t)}^{m \rightarrow m-1} . \\
 &\quad - \overbrace{r p_P(m, t)}^{m \rightarrow m+1, A \rightarrow U} .
 \end{aligned} \tag{A.7}$$

k_P^+ and k_P^- are defined in close analogy to k_R^+ and k_R^- , except for RNAP binding and unbinding instead of repressor. Similarly $p_P(m, t)$ is defined for the active (RNAP-bound) state exactly as are $p_R(m, t)$ and $p_E(m, t)$ for the repressor bound and unbound states, respectively. The new subtlety Eq. A.7 introduces compared to Eq. A.6 is that when mRNAs are produced, the promoter state also changes. This is encoded by the terms involving r , the last term in each of the equations for p_E and p_P . The former accounts for arrivals in the unbound state and the latter accounts for departures from the RNAP-bound state.

To condense and clarify the unwieldy notation of Eq. A.7, it can be rewritten in matrix form. We define the column vector $\vec{p}(m, t)$ as

$$\vec{p}(m, t) = \begin{pmatrix} p_R(m, t) \\ p_E(m, t) \\ p_P(m, t) \end{pmatrix} \tag{A.8}$$

to gather, for a given m , the probabilities of finding the system in the three possible promoter states. Then all the transition rates may be condensed into matrices which multiply this vector. The first matrix is

$$\mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- \\ 0 & k_P^+ & -k_P^- \end{pmatrix}, \quad (\text{A.9})$$

which tracks all promoter state changes in Eq. A.7 that are *not* accompanied by a change in the mRNA copy number. The two terms accounting for transcription, the only transition that increases mRNA copy number, must be handled by two separate matrices given by

$$\mathbf{R}_A = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & r \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & r \end{pmatrix}. \quad (\text{A.10})$$

\mathbf{R}_A accounts for transitions *arriving* in a given state while \mathbf{R}_D tracks *departing* transitions. With these definitions, we can condense Eq. A.7 into the single equation

$$\frac{d}{dt} \vec{p}(m, t) = (\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m, t) + \mathbf{R}_A \vec{p}(m-1, t) + \gamma(m+1) \mathbf{I} \vec{p}(m+1, t), \quad (\text{A.11})$$

which is just Eq. 2.15 in the main text. Straightforward albeit tedious algebra verifies that Eqs. A.7 and A.11 are in fact equivalent.

Although we derived Eq. A.11 for the particular case of nonequilibrium model 2 in Figure 2.1, in fact the chemical master equations for all of the nonequilibrium models in Figure 2.1 except for model 5 can be cast in this form. (We treat model 5 separately in Appendix A.2.) Model 3 introduces no new subtleties beyond model 2 and Eq. A.11 applies equally well to it, simply with different matrices of dimension four instead of three. Models 1 and 4 are both handled by Eq. 2.12 in the main text, which is just Eq. A.11 except in the special case of $\mathbf{R}_D = \mathbf{R}_A \equiv \mathbf{R}$, since in these two models transcription initiation events do not change promoter state.

Recalling that our goal in this section is to derive expressions for mean mRNA and Fano factor for nonequilibrium models 1 through four in Figure 2.1, and since all four of these models are described by Eq. A.11, we can save substantial effort by deriving general formulas for mean mRNA and Fano factor from Eq. A.11 once and for all. Then for each model we can simply plug in the appropriate matrices for \mathbf{K} , \mathbf{R}_D , and \mathbf{R}_A and carry out the remaining algebra.

For our purposes it will suffice to derive the first and second moments of the mRNA distribution from this master equation, similar to the treatment in [1], but we refer the interested reader to [2] for an analogous treatment demonstrating an analytical solution for arbitrary moments.

General forms for mean mRNA and Fano factor

Our task now is to derive expressions for the first two moments of the steady-state mRNA distribution from Eq. A.11. Our treatment of this is analogous to that given in Refs. [1] and [2]. We first obtain the steady-state limit of Eq. A.11 in which the time derivative vanishes, giving

$$0 = (\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m) + \mathbf{R}_A \vec{p}(m-1) + \gamma(m+1) \mathbf{I} \vec{p}(m+1), \quad (\text{A.12})$$

From this, we want to compute

$$\langle m \rangle = \sum_S \sum_{m=0}^{\infty} m p_S(m) \quad (\text{A.13})$$

and

$$\langle m^2 \rangle = \sum_S \sum_{m=0}^{\infty} m^2 p_S(m) \quad (\text{A.14})$$

which define the average values of m and m^2 at steady state, where the averaging is over all possible mRNA copy numbers and promoter states S . For example, for model 1 in Figure 2.1(C), the sum on S would cover repressor bound and unbound states (R and U respectively), for model 2, the sum would cover repressor bound, polymerase bound, and empty states (R , P , and E), and so on for the other models.

Along the way it will be convenient to define the following *conditional* moments as

$$\langle \vec{m} \rangle = \sum_{m=0}^{\infty} m \vec{p}(m), \quad (\text{A.15})$$

and

$$\langle \vec{m}^2 \rangle = \sum_{m=0}^{\infty} m^2 \vec{p}(m). \quad (\text{A.16})$$

These objects are vectors of the same size as $\vec{p}(m)$, and each component can be thought of as the expected value of the mRNA copy number, or copy number squared, conditional on the promoter being in a certain state. For example, for model 1 in Figure 2.1 which has repressor bound and unbound states labeled R and U , $\langle \vec{m}^2 \rangle$ would be

$$\langle \vec{m}^2 \rangle = \begin{pmatrix} \sum_{m=0}^{\infty} m^2 p_R(m) \\ \sum_{m=0}^{\infty} m^2 p_U(m) \end{pmatrix}. \quad (\text{A.17})$$

Analogously to $\langle \vec{m} \rangle$ and $\langle \vec{m}^2 \rangle$, it is convenient to define the vector

$$\langle \vec{m}^0 \rangle = \sum_{m=0}^{\infty} \vec{p}(m), \quad (\text{A.18})$$

whose elements are simply the probabilities of finding the system in each of the possible promoter states. It will be convenient to denote by $\vec{1}^\dagger$ a row vector of the same length as \vec{p} whose elements are all 1, such that, for instance, $\vec{1}^\dagger \cdot \langle \vec{m}^0 \rangle = 1$, $\vec{1}^\dagger \cdot \langle \vec{m} \rangle = \langle m \rangle$, etc.

Promoter state probabilities $\langle \vec{m}^0 \rangle$

To begin, we will find the promoter state probabilities $\langle \vec{m}^0 \rangle$ from Eq. A.12 by summing over all mRNA copy numbers m , producing

$$0 = \sum_{m=0}^{\infty} [(\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m) + \mathbf{R}_A \vec{p}(m-1) + \gamma(m+1) \mathbf{I} \vec{p}(m+1)] \quad (\text{A.19})$$

Using the definitions of $\langle \vec{m}^0 \rangle$ and $\langle \vec{m} \rangle$, and noting the matrices \mathbf{K} , \mathbf{R}_D , and \mathbf{R}_A are all independent of m and can be moved outside the sum, this simplifies to

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m}^0 \rangle - \gamma \langle \vec{m} \rangle + \mathbf{R}_A \sum_{m=0}^{\infty} \vec{p}(m-1) + \gamma \sum_{m=0}^{\infty} (m+1) \vec{p}(m+1). \quad (\text{A.20})$$

The last two terms can be handled by reindexing the summations, transforming them to match the definitions of $\langle \vec{m}^0 \rangle$ and $\langle \vec{m} \rangle$. For the first, noting $\vec{p}(-1) = 0$ since negative numbers of mRNA are nonsensical, we have

$$\sum_{m=0}^{\infty} \vec{p}(m-1) = \sum_{m=-1}^{\infty} \vec{p}(m) = \sum_{m=0}^{\infty} \vec{p}(m) = \langle \vec{m}^0 \rangle. \quad (\text{A.21})$$

Similarly for the second,

$$\sum_{m=0}^{\infty} (m+1) \vec{p}(m+1) = \sum_{m=1}^{\infty} m \vec{p}(m) = \sum_{m=0}^{\infty} m \vec{p}(m) = \langle \vec{m} \rangle, \quad (\text{A.22})$$

which holds since in extending the lower limit from $m = 1$ to $m = 0$, the extra term we added to the sum is zero. Substituting these back in Eq. A.20, we have

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m}^0 \rangle - \gamma \langle \vec{m} \rangle + \mathbf{R}_A \langle \vec{m}^0 \rangle + \gamma \langle \vec{m} \rangle, \quad (\text{A.23})$$

or simply

$$0 = (\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A) \langle \vec{m}^0 \rangle. \quad (\text{A.24})$$

So given matrices \mathbf{K} , \mathbf{R}_D , and \mathbf{R}_A describing a promoter, finding $\langle \vec{m}^0 \rangle$ simply amounts to solving this set of linear equations, subject to the normalization constraint $\vec{1}^\dagger \cdot \langle \vec{m}^0 \rangle = 1$. It will turn out to be the case that, if the matrix $\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A$ is n dimensional, it will always have only $n - 1$ linearly independent equations. Including the normalization condition provides the n -th linearly independent equation, ensuring a unique solution. So when using this equation to solve for $\langle \vec{m}^0 \rangle$, we may always drop one row of the matrix equation at our convenience and supplement the system with the normalization condition instead. The reader may find it illuminating to skip ahead and see Eq. A.24 in use with concrete examples, e.g., Eq. A.52 and what follows it, before continuing on through the general formulas for moments.

First moments $\langle \vec{m} \rangle$ and $\langle m \rangle$

By analogy to the above procedure for finding $\langle \vec{m}^0 \rangle$, we may find $\langle \vec{m} \rangle$ by first multiplying Eq. A.12 by m and then summing over m . Carrying out this procedure we have

$$0 = \sum_{m=0}^{\infty} m [(\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m) + \mathbf{R}_A \vec{p}(m-1) + \gamma(m+1) \mathbf{I} \vec{p}(m+1)], \quad (\text{A.25})$$

and now identifying $\langle \vec{m} \rangle$ and $\langle \vec{m}^2 \rangle$ gives

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m} \rangle - \gamma \langle \vec{m}^2 \rangle + \mathbf{R}_A \sum_{m=0}^{\infty} m \vec{p}(m-1) + \gamma \sum_{m=0}^{\infty} m(m+1) \vec{p}(m+1). \quad (\text{A.26})$$

The summations in the last two terms can be reindexed just as we did for $\langle \vec{m}^0 \rangle$, freely adding or removing terms from the sum which are zero. For the first term we find

$$\sum_{m=0}^{\infty} m \vec{p}(m-1) = \sum_{m=-1}^{\infty} (m+1) \vec{p}(m) = \sum_{m=0}^{\infty} (m+1) \vec{p}(m) = \langle \vec{m} \rangle + \langle \vec{m}^0 \rangle, \quad (\text{A.27})$$

and similarly for the second,

$$\sum_{m=0}^{\infty} m(m+1) \vec{p}(m+1) = \sum_{m=1}^{\infty} (m-1)m \vec{p}(m) = \sum_{m=0}^{\infty} (m-1)m \vec{p}(m) = \langle \vec{m}^2 \rangle - \langle \vec{m} \rangle. \quad (\text{A.28})$$

Substituting back in Eq. A.26 then produces

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m} \rangle - \gamma \langle \vec{m}^2 \rangle + \mathbf{R}_A (\langle \vec{m} \rangle + \langle \vec{m}^0 \rangle) + \gamma (\langle \vec{m}^2 \rangle - \langle \vec{m} \rangle), \quad (\text{A.29})$$

or after simplifying

$$0 = (\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - \gamma) \langle \vec{m} \rangle + \mathbf{R}_A \langle \vec{m}^0 \rangle. \quad (\text{A.30})$$

So like $\langle \vec{m}^0 \rangle$, $\langle \vec{m} \rangle$ is also found by simply solving a set of linear equations after first solving for $\langle \vec{m}^0 \rangle$ from Eq. A.24.

Next we can find the mean mRNA copy number $\langle m \rangle$ from $\langle \vec{m} \rangle$ according to

$$\langle m \rangle = \vec{1}^\dagger \cdot \langle \vec{m} \rangle, \quad (\text{A.31})$$

where $\vec{1}^\dagger$ is a row vector whose elements are all 1. Eq. A.31 holds since the i^{th} element of the column vector $\langle \vec{m} \rangle$ is the mean mRNA value conditional on the system occupying the i^{th} promoter state, so the dot product with $\vec{1}^\dagger$ amounts to simply summing the elements of $\langle \vec{m} \rangle$. Rather than solving Eq. A.30 for $\langle \vec{m} \rangle$ and then computing $\vec{1}^\dagger \cdot \langle \vec{m} \rangle$, we may take a shortcut by multiplying Eq. A.30 by $\vec{1}^\dagger$ first. The key observation that makes this useful is that

$$\vec{1}^\dagger \cdot (\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A) = 0. \quad (\text{A.32})$$

Intuitively, this equality holds because each column of this matrix represents transitions to and from a given promoter state. In any given column, the diagonal encodes all departing transitions and off-diagonals encode all arriving transitions. Conservation of probability means that each column sums to zero, and summing columns is exactly the operation that multiplying by $\vec{1}^\dagger$ carries out.

Proceeding then in multiplying Eq. A.30 by $\vec{1}^\dagger$ produces

$$0 = -\gamma \vec{1}^\dagger \cdot \langle \vec{m} \rangle + \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle, \quad (\text{A.33})$$

or simply

$$\langle m \rangle = \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle. \quad (\text{A.34})$$

We note that the in equilibrium models of transcription such as in Figure 2.1, it is usually *assumed* that the mean mRNA level is given by the ratio of initiation rate r to degradation rate γ multiplied by the probability of finding the system in the RNAP-bound state. Reassuringly, we have recovered exactly this result from the master equation picture: the product $\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle$ picks out the probability of the active promoter state from $\langle \vec{m}^0 \rangle$ and multiplies it by the initiation rate r .

Second moment $\langle m^2 \rangle$ and Fano factor ν

Continuing the pattern of the zeroth and first moments, we now find $\langle \vec{m}^2 \rangle$ by multiplying Eq. A.12 by m^2 and then summing over m , which explicitly is

$$0 = \sum_{m=0}^{\infty} m^2 [(\mathbf{K} - \mathbf{R}_D - \gamma m \mathbf{I}) \vec{p}(m) + \mathbf{R}_A \vec{p}(m-1) + \gamma(m+1) \mathbf{I} \vec{p}(m+1)]. \quad (\text{A.35})$$

Identifying the moments $\langle \vec{m}^2 \rangle$ and $\langle \vec{m}^3 \rangle$ in the first term simplifies this to

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m}^2 \rangle - \gamma \langle \vec{m}^3 \rangle + \mathbf{R}_A \sum_{m=0}^{\infty} m^2 \vec{p}(m-1) + \gamma \sum_{m=0}^{\infty} m^2 (m+1) \vec{p}(m+1). \quad (\text{A.36})$$

Reindexing the sums of the last two terms proceeds just as it did for the zeroth and first moments. Explicitly, we have

$$\sum_{m=0}^{\infty} m^2 \vec{p}(m-1) = \sum_{m=-1}^{\infty} (m+1)^2 \vec{p}(m) = \sum_{m=0}^{\infty} (m+1)^2 \vec{p}(m) = \langle \vec{m}^2 \rangle + 2\langle \vec{m} \rangle + \langle \vec{m}^0 \rangle, \quad (\text{A.37})$$

for the first sum and

$$\sum_{m=0}^{\infty} m^2 (m+1) \vec{p}(m+1) = \sum_{m=1}^{\infty} (m-1)^2 m \vec{p}(m) = \sum_{m=0}^{\infty} (m-1)^2 m \vec{p}(m) = \langle \vec{m}^3 \rangle - 2\langle \vec{m}^2 \rangle + \langle \vec{m} \rangle \quad (\text{A.38})$$

for the second. Substituting the results of the sums back in Eq. A.36 gives

$$0 = (\mathbf{K} - \mathbf{R}_D) \langle \vec{m}^2 \rangle - \gamma \langle \vec{m}^3 \rangle + \mathbf{R}_A (\langle \vec{m}^2 \rangle + 2\langle \vec{m} \rangle + \langle \vec{m}^0 \rangle) + \gamma (\langle \vec{m}^3 \rangle - 2\langle \vec{m}^2 \rangle + \langle \vec{m} \rangle), \quad (\text{A.39})$$

and after grouping like powers of m we have

$$0 = (\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - 2\gamma) \langle \vec{m}^2 \rangle + (2\mathbf{R}_A + \gamma) \langle \vec{m} \rangle + \mathbf{R}_A \langle \vec{m}^0 \rangle. \quad (\text{A.40})$$

As we found when computing $\langle m \rangle$ from $\langle \vec{m} \rangle$, we can spare ourselves some algebra by multiplying Eq. A.40 by $\vec{1}^\dagger$, which then reduces to

$$0 = -2\gamma \langle m^2 \rangle + \vec{1}^\dagger \cdot (2\mathbf{R}_A + \gamma) \langle \vec{m} \rangle + \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle, \quad (\text{A.41})$$

and noting from Eq. A.34 that $\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle = \gamma \langle m \rangle$, we have the tidy result

$$\langle m^2 \rangle = \langle m \rangle + \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle. \quad (\text{A.42})$$

Finally we have all the preliminary results needed to write a general expression for the Fano factor ν . The Fano factor is defined as the ratio of variance to mean, which can be written as

$$\nu = \frac{\langle m^2 \rangle - \langle m \rangle^2}{\langle m \rangle} = \frac{\langle m \rangle + \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle - \langle m \rangle^2}{\langle m \rangle} \quad (\text{A.43})$$

and simplified to

$$\nu = 1 - \langle m \rangle + \frac{\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle}{\gamma \langle m \rangle}. \quad (\text{A.44})$$

Note a subtle notational trap here: $\langle m \rangle = \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle$ rather than the by-eye similar but wrong expression $\langle m \rangle \neq \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle$, so the last term in Eq. A.44 is in general quite nontrivial. For a generic promoter, Eq. A.44 may be greater than, less than, or equal to one, as asserted in Section 2.4. We have not found the general form Eq. A.44 terribly intuitive and instead defer discussion to specific examples.

Summary of general results

For ease of reference, we collect and reprint here the key results derived in this section that are used in the main text and subsequent subsections. Mean mRNA copy number and Fano factor are given by Eqs. A.34 and A.44, which are

$$\langle m \rangle = \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle \quad (\text{A.45})$$

and

$$\nu = 1 - \langle m \rangle + \frac{\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle}{\gamma \langle m \rangle}, \quad (\text{A.46})$$

respectively. To compute these two quantities, we need the expressions for $\langle \vec{m}^0 \rangle$ and $\langle \vec{m} \rangle$ given by solving Eqs. A.24 and A.30, respectively, which are

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A) \langle \vec{m}^0 \rangle = 0 \quad (\text{A.47})$$

and

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - \gamma \mathbf{I}) \langle \vec{m} \rangle = -\mathbf{R}_A \langle \vec{m}^0 \rangle. \quad (\text{A.48})$$

Some comments are in order before we consider particular models. First, note that to obtain $\langle \vec{m} \rangle$ and ν , we need not bother solving for all components of the vectors $\langle \vec{m}^0 \rangle$ and $\langle \vec{m} \rangle$, but only the components which are multiplied by nonzero elements of \mathbf{R}_A . The only component of $\langle \vec{m}^0 \rangle$ that ever survives is the transcriptionally active state, and for the models we consider here, there is only ever one such state. This will save us some amount of algebra below.

Also note that we are computing Fano factors to verify the results of Section 2.4, concerning the constitutive promoter models in Figure 2.2 which are analogs of the simple repression models in Figure 2.1. We can translate the matrices from the simple repression models to the constitutive case by simply substituting all occurrences of repressor rates by zero and removing the row and column corresponding to the repressor bound state. The results for $\langle m \rangle$ computed in the repressed case can be easily translated to the constitutive case, rather than recalculating from scratch, by

taking the limit $k_R^+ \rightarrow 0$, since this amounts to sending repressor copy number to zero.

Finally, we point out that it would be possible to compute $\langle \vec{m}^0 \rangle$ more simply using the diagram methods from King and Altman [3] (also independently discovered by Hill [4]). But to our knowledge this method cannot be applied to compute $\langle \vec{m} \rangle$ or ν , so since we would need to resort to solving the matrix equations anyways for $\langle \vec{m} \rangle$, we choose not to introduce the extra conceptual burden of the diagram methods simply for computing $\langle \vec{m}^0 \rangle$.

Nonequilibrium Model One - Poisson Promoter

Mean mRNA

For nonequilibrium model 1 in Figure 2.1, we have already shown the full master equation in Eq. 2.10 and Eq. A.6, but for completeness we reprint it again as

$$\begin{aligned} \frac{d}{dt} p_R(m, t) &= - \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} + \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_R(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_R(m, t)}^{m \rightarrow m-1} \\ \frac{d}{dt} p_U(m, t) &= \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} - \overbrace{k_R^+ p_U(m, t)}^{U \rightarrow R} + \overbrace{r p_U(m-1, t)}^{m-1 \rightarrow m} - \overbrace{r p_U(m, t)}^{m \rightarrow m+1} \\ &\quad + \overbrace{(m+1)\gamma p_U(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_U(m, t)}^{m \rightarrow m-1}. \end{aligned} \tag{A.49}$$

This is a direct transcription of the states and rates in Figure 2.1. This may be converted to the matrix form of the master equation shown in Eq. A.11 with matrices

$$\vec{p}(m) = \begin{pmatrix} p_R(m) \\ p_U(m) \end{pmatrix}, \quad \mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ \\ k_R^- & -k_R^+ \end{pmatrix}, \quad \mathbf{R} = \begin{pmatrix} 0 & 0 \\ 0 & r \end{pmatrix}, \tag{A.50}$$

where \mathbf{R}_A and \mathbf{R}_D are equal, so we drop the subscript and denote both simply by \mathbf{R} . Since our interest is only in steady-state we dropped the time dependence as well.

First we need $\langle \vec{m}^0 \rangle$. Label its components as p_R and p_U , the probabilities of finding the system in either promoter state, and note that only p_U survives multiplication by \mathbf{R} , since

$$\mathbf{R} \langle \vec{m}^0 \rangle = \begin{pmatrix} 0 & 0 \\ 0 & r \end{pmatrix} \begin{pmatrix} p_R \\ p_U \end{pmatrix} = \begin{pmatrix} 0 \\ r p_U \end{pmatrix}, \tag{A.51}$$

so we need not bother finding p_R . Then we have

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A) \langle \vec{m}^0 \rangle = \begin{pmatrix} -k_R^- & k_R^+ \\ k_R^- & -k_R^+ \end{pmatrix} \begin{pmatrix} p_R \\ p_U \end{pmatrix} = 0. \tag{A.52}$$

As mentioned earlier in Section A.1, the two rows are linearly dependent, so taking only the first row and using normalization to set $p_R = 1 - p_U$ gives

$$-k_R^-(1 - p_U) + k_R^+ p_U = 0, \quad (\text{A.53})$$

which is easily solved to find

$$p_U = \frac{k_R^-}{k_R^- + k_R^+}. \quad (\text{A.54})$$

Substituting this into Eq. A.51, and the result of that into Eq. A.45, we have

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_R^-}{k_R^- + k_R^+} \quad (\text{A.55})$$

as asserted in Eq. 2.13 of the main text.

Fano factor

To verify that the Fano factor for model 1 in Figure 2.2(A) is in fact 1 as claimed in the main text, note that in this limit $p_U = 1$ and $\langle m \rangle = r/\gamma$. All elements of \mathbf{K} are zero, and $\mathbf{R}_A - \mathbf{R}_D = 0$, so Eq. A.48 reduces to

$$-\gamma \langle \vec{m} \rangle = -r, \quad (\text{A.56})$$

or, in other words, since there is only one promoter state, $\langle \vec{m} \rangle = \langle m \rangle$. Then it follows that

$$\nu = 1 - \frac{r}{\gamma} + \frac{r \langle m \rangle}{\gamma \langle m \rangle} = 1 \quad (\text{A.57})$$

as claimed.

Nonequilibrium Model Two - RNAP Bound and Unbound States Mean mRNA

As shown earlier, the full master equation for model 2 in Figure 2.1 is

$$\begin{aligned}
 \frac{d}{dt}p_R(m, t) &= - \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} + \overbrace{k_R^+ p_E(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_R(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_R(m, t)}^{m \rightarrow m-1} \\
 \frac{d}{dt}p_E(m, t) &= \overbrace{k_R^- p_R(m, t)}^{R \rightarrow U} - \overbrace{k_R^+ p_E(m, t)}^{U \rightarrow R} + \overbrace{(m+1)\gamma p_E(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_E(m, t)}^{m \rightarrow m-1} . \\
 &\quad + \overbrace{k_P^- p_P(m, t)}^{A \rightarrow U} - \overbrace{k_P^+ p_E(m, t)}^{U \rightarrow A} + \overbrace{r p_P(m-1, t)}^{m-1 \rightarrow m, A \rightarrow U} \\
 \frac{d}{dt}p_P(m, t) &= - \overbrace{k_P^- p_P(m, t)}^{A \rightarrow U} + \overbrace{k_P^+ p_E(m, t)}^{U \rightarrow A} + \overbrace{(m+1)\gamma p_P(m+1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma m p_P(m, t)}^{m \rightarrow m-1} . \\
 &\quad - \overbrace{r p_P(m, t)}^{m \rightarrow m+1, A \rightarrow U} ,
 \end{aligned} \tag{A.58}$$

which can be condensed to the matrix form of Eq. A.11 with matrices given by

$$\mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- \\ 0 & k_P^+ & -k_P^- \end{pmatrix}, \quad \mathbf{R}_A = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & r \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & r \end{pmatrix}. \tag{A.59}$$

As for model 1, we must first find $\mathbf{R}_A \langle \vec{m}^0 \rangle$. Denote its components as p_R , p_E , p_P , the probabilities of being found in repressor bound, empty, or RNAP-bound states, respectively. Only p_P is necessary to find since

$$\mathbf{R}_A \langle \vec{m}^0 \rangle = \begin{pmatrix} 0 \\ r p_P \\ 0 \end{pmatrix}. \tag{A.60}$$

Then Eq. A.47 for $\langle \vec{m} \rangle$ reads

$$\begin{pmatrix} -k_R^- & k_R^+ & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- + r \\ 0 & k_P^+ & -k_P^- - r \end{pmatrix} \begin{pmatrix} p_R \\ p_E \\ p_P \end{pmatrix} = 0. \tag{A.61}$$

Discarding the middle row as redundant and incorporating the normalization condition leads to a set of three linearly independent equations, namely

$$-k_R^- p_R + k_R^+ p_E = 0 \quad (\text{A.62})$$

$$k_P^+ p_E + (-k_P^- - r) p_P = 0 \quad (\text{A.63})$$

$$p_R + p_E + p_P = 1. \quad (\text{A.64})$$

Using $p_R = 1 - p_E - p_P$ to eliminate p_R in the first and solving the resulting equation for p_E gives $p_E = (1 - p_P)k_R^- / (k_R^- + k_R^+)$. Substituting this for p_E in the second equation gives an equation in p_P alone which is

$$k_P^+ k_R^- (1 - p_P) - (k_P^- + r)(k_R^+ + k_R^-) p_P = 0 \quad (\text{A.65})$$

and solving for p_P gives

$$p_P = \frac{k_P^+ k_R^-}{k_P^+ k_R^- + (k_P^- + r)(k_R^+ + k_R^-)}. \quad (\text{A.66})$$

Substituting this in Eq. A.60 and multiplying by \mathbf{R}_A produces

$$\mathbf{R}_A \langle \vec{m}^0 \rangle = r \frac{k_P^+ k_R^-}{k_P^+ k_R^- + (k_P^- + r)(k_R^+ + k_R^-)} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \quad (\text{A.67})$$

from which $\langle m \rangle$ follows readily,

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_P^+ k_R^-}{k_P^+ k_R^- + (k_P^- + r)(k_R^+ + k_R^-)}, \quad (\text{A.68})$$

as claimed in Eq. 2.18 in the main text.

Fano factor

To compute the Fano factor, we first remove the repressor bound state from the matrices describing the model, which reduce to

$$\mathbf{K} = \begin{pmatrix} -k_P^+ & k_P^- \\ k_P^+ & -k_P^- \end{pmatrix}, \quad \mathbf{R}_A = \begin{pmatrix} 0 & r \\ 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 \\ 0 & r \end{pmatrix}. \quad (\text{A.69})$$

Similarly we remove the repressor bound state from $\mathbf{R}_A \langle \vec{m}^0 \rangle$ and take the $k_R^+ \rightarrow 0$ limit, which simplifies to

$$\mathbf{R}_A \langle \vec{m}^0 \rangle = r \frac{k_P^+}{k_P^+ + k_P^- + r} \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \quad (\text{A.70})$$

Then we must compute $\langle \vec{m} \rangle$ from Eq. A.48, which with these matrices reads

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - \gamma \mathbf{I}) \langle \vec{m} \rangle = \begin{pmatrix} -k_P^+ - \gamma & k_P^- + r \\ k_P^+ & -k_P^- - r - \gamma \end{pmatrix} \begin{pmatrix} m_E \\ m_P \end{pmatrix} = r \frac{k_P^+}{k_P^+ + k_P^- + r} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad (\text{A.71})$$

where we labeled the components of $\langle \vec{m} \rangle$ as m_E and m_P , since they are the mean mRNA counts conditional upon the system residing in the empty or polymerase bound states, respectively. Unlike for $\langle \vec{m}^0 \rangle$, the rows of this matrix are linearly independent so we simply solve this matrix equation as is. We can immediately eliminate m_E since $m_E = m_P(k_P^- + r + \gamma)/k_P^+$ from the second row, and substituting into the first row gives an equation for m_P alone, which is

$$\left[-(k_P^+ + \gamma)(k_P^- + r + \gamma) + k_P^+(k_P^- + r) \right] m_P = -\frac{r(k_P^+)^2}{k_P^+ + k_P^- + r}. \quad (\text{A.72})$$

Expanding the products cancels several terms, and solving for m_P gives

$$m_P = \frac{r(k_P^+)^2}{\gamma(k_P^+ + k_P^- + r)(k_P^+ + k_P^- + r + \gamma)}. \quad (\text{A.73})$$

Note then that $\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle = r m_P$. We also need the constitutive limit of $\langle m \rangle$ from Eq. A.68, again found by taking $k_R^+ \rightarrow 0$, which is

$$\langle m \rangle = \frac{r}{\gamma} \frac{k_P^+}{k_P^+ + k_P^- + r} \quad (\text{A.74})$$

and substituting this along with $\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle = r m_P$ into Eq. A.46 for the Fano factor ν , we find

$$\nu = 1 - \frac{r}{\gamma} \frac{k_P^+}{k_P^+ + k_P^- + r} + \frac{r}{\gamma} \frac{r(k_P^+)^2}{\gamma(k_P^+ + k_P^- + r)(k_P^+ + k_P^- + r + \gamma)} \left(\frac{r}{\gamma} \frac{k_P^+}{k_P^+ + k_P^- + r} \right)^{-1}. \quad (\text{A.75})$$

This simplifies to

$$\nu = 1 - \frac{r}{\gamma} \left(\frac{k_P^+}{k_P^+ + k_P^- + r} - \frac{k_P^+}{k_P^+ + k_P^- + r + \gamma} \right), \quad (\text{A.76})$$

which further simplifies to

$$\nu = 1 - \frac{r k_P^+}{(k_P^+ + k_P^- + r)(k_P^+ + k_P^- + r + \gamma)}, \quad (\text{A.77})$$

exactly Eq. 2.36 in the main text.

Nonequilibrium Model Three - Multistep Transcription Initiation and Escape Mean mRNA

In close analogy to model 2 above, nonequilibrium model 3 from Figure 2.1(C) can be described by our generic master equation Eq. A.11 with promoter transition matrix given by

$$\mathbf{K} = \begin{pmatrix} -k_R^- & k_R^+ & 0 & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- & 0 \\ 0 & k_P^+ & -k_P^- - k_O & 0 \\ 0 & 0 & k_O & 0 \end{pmatrix} \quad (\text{A.78})$$

and transcription matrices given by

$$\mathbf{R}_A = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{R}_D = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r \end{pmatrix}. \quad (\text{A.79})$$

$\langle \vec{m}^0 \rangle$ is again given by Eq. A.47, which in this case takes the form

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A) \langle \vec{m}^0 \rangle = \begin{pmatrix} -k_R^- & k_R^+ & 0 & 0 \\ k_R^- & -k_R^+ - k_P^+ & k_P^- & r \\ 0 & k_P^+ & -k_P^- - k_O & 0 \\ 0 & 0 & k_O & -r \end{pmatrix} \begin{pmatrix} p_R \\ p_E \\ p_C \\ p_O \end{pmatrix} = 0, \quad (\text{A.80})$$

where the four components of $\langle \vec{m}^0 \rangle$ correspond to the four promoter states repressor bound, empty, RNAP-bound closed complex, and RNAP-bound open complex. As explained in Section A.1, we are free to discard one linearly dependent row from this matrix and replace it with the normalization condition $p_R + p_E + p_C + p_O = 1$. Using normalization to eliminate p_R from the first row gives

$$p_E = (1 - p_C - p_O) \frac{k_R^-}{k_R^- + k_R^+}. \quad (\text{A.81})$$

If we substitute this in the third row, then the last two rows constitute two equations in p_C and p_O given by

$$k_P^+ k_R^- (1 - p_C - p_O) - (k_P^- + k_O) (k_R^+ + k_R^-) p_C = 0 \quad (\text{A.82})$$

$$k_O p_C - r p_O = 0. \quad (\text{A.83})$$

Solving for $p_C = p_O r / k_O$ in the second and substituting into the first gives us our desired single equation in the single variable p_O , which is

$$k_P^+ k_R^- - k_P^+ k_R^- \left(1 + \frac{r}{k_O} \right) p_O - (k_P^- + k_O) (k_R^+ + k_R^-) \frac{r}{k_O} p_O = 0, \quad (\text{A.84})$$

and solving for p_O we find

$$p_O = \frac{k_P^+ k_R^- k_O}{k_P^+ k_R^- k_O + r k_P^+ k_R^- + r(k_P^- + k_O)(k_R^+ + k_R^-)}. \quad (\text{A.85})$$

Once again p_O , the transcriptionally active state, is the only component of $\langle \vec{m}^0 \rangle$ that survives multiplication by \mathbf{R}_A , and $\mathbf{R}_A \langle \vec{m}^0 \rangle = r p_O$. So

$$\langle m \rangle = \frac{1}{\gamma} \vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m}^0 \rangle = \frac{r}{\gamma} \frac{k_P^+ k_R^- k_O}{k_P^+ k_R^- k_O + r k_P^+ k_R^- + r(k_P^- + k_O)(k_R^+ + k_R^-)}, \quad (\text{A.86})$$

which equals Eq. 2.25 in the main text.

Fano factor

To compute the Fano factor of the analogous constitutive promoter, we first excise the repressor states and rates from the problem. More precisely, we construct the matrix $(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - \gamma \mathbf{I})$ and substitute it into Eq. A.48 which is now

$$(\mathbf{K} - \mathbf{R}_D + \mathbf{R}_A - \gamma \mathbf{I}) \langle \vec{m} \rangle = \begin{pmatrix} -k_P^+ - \gamma & k_P^- & r \\ k_P^+ & -k_P^- - k_O - \gamma & 0 \\ 0 & k_O & -r - \gamma \end{pmatrix} \begin{pmatrix} m_E \\ m_C \\ m_O \end{pmatrix} = -r p_O \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \quad (\text{A.87})$$

where we labeled the unbound, closed complex, and open complex components of $\langle \vec{m} \rangle$ as m_E , m_C , and m_O , respectively. p_O is given by the limit of Eq. A.85 as $k_R^+ \rightarrow 0$, which is

$$p_O = \frac{k_P^+ k_O}{k_P^+ (k_O + r) + r(k_P^- + k_O)} \equiv \frac{k_P^+ k_O}{\mathcal{Z}}, \quad (\text{A.88})$$

where we define \mathcal{Z} for upcoming convenience as this sum of terms will appear repeatedly. We can use the second equation to eliminate m_E , finding $m_E = m_C(k_P^- + k_O + \gamma)/k_P^+$, and the third to eliminate m_C , which is simply $m_C = m_O(r + \gamma)/k_O$. Substituting these both into the first equation gives a single equation for the variable of interest, m_O ,

$$-(k_P^+ + \gamma)(k_P^- + k_O + \gamma)(r + \gamma)m_O + k_P^- k_P^+ (r + \gamma)m_O + r k_P^+ k_O m_O = -r k_P^+ k_O p_O, \quad (\text{A.89})$$

which is solved for m_O to give

$$m_O = p_O \frac{r k_P^+ k_O}{(k_P^+ + \gamma)(k_P^- + k_O + \gamma)(r + \gamma) - r k_P^+ k_O - k_P^- k_P^+ (r + \gamma)}. \quad (\text{A.90})$$

Expanding the denominator and canceling terms leads to

$$m_O = p_O \frac{r}{\gamma} \frac{k_P^+ k_O}{\mathcal{Z} + \gamma(k_P^+ + k_P^- + k_O + r) + \gamma^2}. \quad (\text{A.91})$$

Now $\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle = r m_O$, and $\langle m \rangle = r p_O / \gamma$, so if we substitute these two quantities into Eq. A.46, we will readily obtain the Fano factor as

$$\nu = 1 - \langle m \rangle + \frac{\vec{1}^\dagger \cdot \mathbf{R}_A \langle \vec{m} \rangle}{\gamma \langle m \rangle} = 1 - \frac{r}{\gamma} p_O + \frac{m_O}{p_O}. \quad (\text{A.92})$$

Substituting, we see that

$$\nu = 1 - \frac{r}{\gamma} \frac{k_P^+ k_O}{\mathcal{Z}} + \frac{r}{\gamma} \frac{k_P^+ k_O}{\mathcal{Z} + \gamma(k_P^+ + k_P^- + k_O + r) + \gamma^2}, \quad (\text{A.93})$$

and after simplifying, we obtain

$$\nu = 1 - \frac{r k_P^+ k_O}{\mathcal{Z}} \frac{k_P^+ + k_P^- + k_O + r + \gamma}{\mathcal{Z} + \gamma(k_P^+ + k_P^- + k_O + r) + \gamma^2}, \quad (\text{A.94})$$

as stated in Eq. 2.37 in the main text.

Generalizing $\nu < 1$ to more fine-grained models

In the main text we argued that the convolution of multiple exponential distributions should be a distribution with a smaller fractional width than the corresponding exponential distribution with the same mean. This can be made more precise with a result from [5], who showed that the convolution of multiple gamma distributions (of which the exponential distribution is a special case) is, to a very good approximation, also gamma distributed. Using their Eq. 2 for the distribution of the convolution, with shape parameters set to 1 to give exponential distributions, the total waiting time distribution has a ratio of variance to squared mean $\sigma^2 / \mu^2 = \sum_i k_i^2 / (\sum_i k_i)^2$, where the k_i are the rates of the individual steps. Clearly this is less than 1 and therefore the total waiting time distribution is narrower than the corresponding exponential.

We also claimed in the main text that for a process whose waiting time distribution is narrower than exponential, i.e., has $\sigma^2 / \mu^2 < 1$, the distribution of counts should be less variable than a Poisson distribution, leading to a Fano factor $\nu < 1$. This we argue by analogy to photon statistics where it is known that “antibunched” arrivals, in other words more uniformly distributed in time relative to uncorrelated arrivals, generally gives rise to sub-Poissonian noise [6, 7]. Although loopholes to this result are known to exist, to our knowledge they appear to arise from uniquely quantum

effects so we do not expect they apply for our problem. Nevertheless we refrain from elevating this equivalence of kinetic cycles with sub-Poissonian noise to a “theorem.”

Nonequilibrium Model Four - “Active” and “Inactive” States Mean mRNA

The mathematical specification of this model is almost identical to model 2. The matrix \mathbf{K} is identical, as is \mathbf{R}_D . The only difference is that now $\mathbf{R}_A = \mathbf{R}_D$, i.e., both are diagonal, in contrast to model 2 where \mathbf{R}_A has an off-diagonal element, as in Eq. A.59. Then the analog of Eq. A.61 for finding $\langle m^0 \rangle$ is

$$\begin{pmatrix} -k_R^- & k_R^+ & 0 \\ k_R^- & -k_R^+ - k^+ & k^- \\ 0 & k^+ & -k^- \end{pmatrix} \begin{pmatrix} p_R \\ p_I \\ p_A \end{pmatrix} = 0. \quad (\text{A.95})$$

In fact we need not do this calculation explicitly and can instead recycle the calculation of mean mRNA $\langle m \rangle$ from model 2. The matrices are identical except for the relabeling $k^- \longleftrightarrow (k_p^- + r)$, and a careful look through the derivation of $\langle m \rangle$ for model 2 shows that the parameters k_p^- and r only ever appear as the sum $k_p^- + r$. So taking $\langle m \rangle$ from model 2, Eq. A.68, and relabeling $(k_p^- + r) \rightarrow k^-$ gives us our answer for model four, simply

$$\langle m \rangle = \frac{r}{\gamma} \frac{k^+ k_R^-}{k^+ k_R^- + k^- (k_R^+ + k_R^-)}. \quad (\text{A.96})$$

Fano factor

Likewise, for computing the Fano factor of this model we may take a shortcut. Consider the constitutive model four from Figure 2.2 for which we want to compute the Fano factor and compare it to nonequilibrium model one of simple repression in Figure 2.1. Mathematically these are exactly the same model, just with rates labeled differently and the meaning of the promoter states interpreted differently. Furthermore, nonequilibrium model 1 from Figure 2.1 was the model considered by Jones et. al. [8], where they derived the Fano factor for that model to be

$$\nu = 1 + \frac{r k_R^+}{(k_R^+ + k_R^-)(k_R^+ + k_R^- + \gamma)}. \quad (\text{A.97})$$

So recognizing that the relabelings $k_R^+ \rightarrow k^-$ and $k_R^- \rightarrow k^+$ will translate this result to our model four from Figure 2.2, we can immediately write down our Fano factor

as

$$\nu = 1 + \frac{rk^-}{(k^- + k^+)(k^- + k^+ + \gamma)}, \quad (\text{A.98})$$

as quoted in Eq. 2.38 and in Figure 2.2.

A.2 Bursty promoter models - generating function solutions and numerics

Constitutive promoter with bursts

From master equation to generating function

The objective of this section is to write down the steady-state mRNA distribution for model 5 in Figure 2.2. Our claim is that this model is rich enough that it can capture the expression pattern of bacterial constitutive promoters. Figure A.2 shows two different schematic representations of the model. Figure A.2(A) shows the promoter cartoon model with burst initiation rate k_i , mRNA degradation rate γ , and mean burst size b . For our derivation of the chemical master equation we will focus more on Figure A.2(B). This representation is intended to highlight that bursty gene expression allows transitions between mRNA count m and m' even with $m - m' > 1$.




Figure A.2: **Bursty transcription for unregulated promoter.** (A) Schematic of the one-state bursty transcription model. Rate k_i is the bursty initiation rate, γ is the mRNA degradation rate, and b is the mean burst size. (B) Schematic depiction of the mRNA count state transitions. The model in (A) allows for transitions of > 1 mRNA counts with probability $G_{m-m'}$, where the state jumps from having m' mRNA to having m mRNA in a single burst of gene expression.

To derive the master equation we begin by considering the possible state transitions to “enter” state m . There are two possible paths to jump from an mRNA count $m' \neq m$ to a state m in a small time window Δt :

1. By degradation of a single mRNA, jumping from $m + 1$ to m .
2. By producing $m - m'$ mRNA for $m' \in \{0, 1, \dots, m - 1\}$.

For the “exit” states from m into $m' \neq m$ during a small time window Δt we also have two possibilities:

1. By degradation of a single mRNA, jumping from m to $m - 1$.
2. By producing $m' - m$ mRNA for $m' - m \in \{1, 2, \dots\}$.

This implies that the probability of having m mRNA at time $t + \Delta t$ can be written as

$$\begin{aligned}
 p(m, t + \Delta t) = & \overbrace{p(m, t) + \gamma \Delta t (m + 1) p(m + 1, t)}^{m+1 \rightarrow m} - \overbrace{\gamma \Delta t m p(m, t)}^{m \rightarrow m-1} \\
 & + \underbrace{k_i \Delta t \sum_{m'=0}^{m-1} G_{m-m'} p(m', t)}_{m' \in \{0, 1, \dots, m-1\} \rightarrow m} - \underbrace{k_i \Delta t \sum_{m'=m+1}^{\infty} G_{m'-m} p(m, t)}_{m \rightarrow m' \in \{m+1, m+2, \dots\}} \quad (\text{A.99})
 \end{aligned}$$

where we indicate $G_{m'-m}$ as the probability of having a burst of size $m' - m$, i.e. when the number of mRNAs jump from m to $m' > m$ due to a single mRNA transcription burst. We suggestively use the letter G as we will assume that these bursts sizes are geometrically distributed with parameter θ . This is written as

$$G_k = \theta(1 - \theta)^k \text{ for } k \in \{0, 1, 2, \dots\}. \quad (\text{A.100})$$

In Section 2.4 of the main text we derive this functional form for the burst size distribution. An intuitive way to think about it is that for transcription initiation events that take place instantaneously there are two competing possibilities: Producing another mRNA with probability $(1 - \theta)$, or ending the burst with probability θ . What this implies is that for a geometrically distributed burst size we have a mean burst size b of the form

$$b \equiv \langle m' - m \rangle = \sum_{k=0}^{\infty} k \theta (1 - \theta)^k = \frac{1 - \theta}{\theta}. \quad (\text{A.101})$$

To clean up Equation A.99 we can send the first term on the right hand side to the left, and divide both sides by Δt . Upon taking the limit where $\Delta t \rightarrow 0$ we can write

$$\frac{d}{dt} p(m, t) = (m+1)\gamma p(m+1, t) - m\gamma p(m, t) + k_i \sum_{m'=0}^{m-1} G_{m-m'} p(m', t) - k_i \sum_{m'=m+1}^{\infty} G_{m'-m} p(m, t). \quad (\text{A.102})$$

Furthermore, given that the timescale for this equation is set by the mRNA degradation rate γ we can divide both sides by this rate, obtaining

$$\frac{d}{d\tau}p(m, \tau) = (m+1)p(m+1, \tau) - mp(m, \tau) + \lambda \sum_{m'=0}^{m-1} G_{m-m'}p(m', \tau) - \lambda \sum_{m'=m+1}^{\infty} G_{m'-m}p(m, \tau), \quad (\text{A.103})$$

where we defined $\tau \equiv t \times \gamma$, and $\lambda \equiv k_i/\gamma$. The last term in Eq. A.103 sums all burst sizes except for a burst of size zero. We can re-index the sum to include this term, obtaining

$$\lambda \sum_{m'=m+1}^{\infty} G_{m'-m}p(m, \tau) = \lambda p(m, \tau) \left[\underbrace{\sum_{m'=m}^{\infty} G_{m'-m}}_{\text{re-index sum to include burst size zero}} - \underbrace{G_0}_{\text{subtract extra added term}} \right]. \quad (\text{A.104})$$

Given the normalization constraint of the geometric distribution, adding the probability of all possible burst sizes – including size zero since we re-indexed the sum – allows us to write

$$\sum_{m'=m}^{\infty} G_{m'-m} - G_0 = 1 - G_0. \quad (\text{A.105})$$

Substituting this into Eq. A.103 results in

$$\frac{d}{d\tau}p(m, \tau) = (m+1)p(m+1, \tau) - mp(m, \tau) + \lambda \sum_{m'=0}^{m-1} G_{m-m'}p(m', \tau) - \lambda p(m, \tau) [1 - G_0]. \quad (\text{A.106})$$

To finally get at a more compact version of the equation notice that the third term in Eq. A.106 includes burst from size $m' - m = 1$ to size $m' - m = m$. We can include the term $p(m, \tau)G_0$ in the sum which allows bursts of size $m' - m = 0$. This results in our final form for the chemical master equation

$$\frac{d}{d\tau}p(m, \tau) = (m+1)p(m+1, \tau) - mp(m, \tau) - \lambda p(m, \tau) + \lambda \sum_{m'=0}^m G_{m-m'}p(m', \tau). \quad (\text{A.107})$$

In order to solve Eq. A.107 we will use the generating function method [9]. The probability generating function is defined as

$$F(z, t) = \sum_{m=0}^{\infty} z^m p(m, t), \quad (\text{A.108})$$

where z is just a dummy variable that will help us later on to obtain the moments of the distribution. Let us now multiply both sides of Eq. A.107 by z^m and sum over all m

$$\sum_m z^m \frac{d}{d\tau} p(m, \tau) = \sum_m z^m \left[-mp(m, \tau) + (m+1)p(m+1, \tau) + \lambda \sum_{m'=0}^m G_{m-m'} p(m', \tau) - \lambda p(m, \tau) \right], \quad (\text{A.109})$$

where we use $\sum_m \equiv \sum_{m=0}^{\infty}$. We can distribute the sum and use the definition of $F(z, \tau)$ to obtain

$$\frac{dF(z, \tau)}{d\tau} = - \sum_m z^m m p(m, \tau) + \sum_m z^m (m+1) p(m+1, \tau) + \lambda \sum_m z^m \sum_{m'=0}^m G_{m-m'} p(m', \tau) - \lambda F(z, \tau). \quad (\text{A.110})$$

We can make use of properties of the generating function to write everything in terms of $F(z, \tau)$: the first term on the right hand side of Eq. A.110 can be rewritten as

$$\sum_m z^m \cdot m \cdot p(m, \tau) = \sum_m z \frac{\partial z^m}{\partial z} p(m, \tau), \quad (\text{A.111})$$

$$= \sum_m z \frac{\partial}{\partial z} (z^m p(m, \tau)), \quad (\text{A.112})$$

$$= z \frac{\partial}{\partial z} \left(\sum_m z^m p(m, \tau) \right), \quad (\text{A.113})$$

$$= z \frac{\partial F(z, \tau)}{\partial z}. \quad (\text{A.114})$$

For the second term on the right hand side of Eq. A.110 we define $k \equiv m+1$. This allows us to write

$$\sum_{m=0}^{\infty} z^m \cdot (m+1) \cdot p(m+1, \tau) = \sum_{k=1}^{\infty} z^{k-1} \cdot k \cdot p(k, \tau), \quad (\text{A.115})$$

$$= z^{-1} \sum_{k=1}^{\infty} z^k \cdot k \cdot p(k, \tau), \quad (\text{A.116})$$

$$= z^{-1} \sum_{k=0}^{\infty} z^k \cdot k \cdot p(k, \tau), \quad (\text{A.117})$$

$$= z^{-1} \left(z \frac{\partial F(z)}{\partial z} \right), \quad (\text{A.118})$$

$$= \frac{\partial F(z)}{\partial z}. \quad (\text{A.119})$$




Figure A.3: **Reindexing double sum.** Schematic for reindexing the sum $\sum_{m=0}^{\infty} \sum_{m'=0}^m$. Blue circles depict the 2D grid of nonnegative integers restricted to the lower triangular part of the m, m' plane. The trick is that this double sum runs over all (m, m') pairs with $m' \leq m$. Summing m first instead of m' requires determining the boundary: the upper boundary of the m' -first double sum becomes the lower boundary of the m -first double sum.

The third term in Eq. A.110 is the most trouble. The trick is to reverse the default order of the sums as

$$\sum_{m=0}^{\infty} \sum_{m'=0}^m = \sum_{m'=0}^{\infty} \sum_{m=m'}^{\infty}. \quad (\text{A.120})$$

To see the logic of the sum we point the reader to Figure A.3. The key is to notice that the double sum $\sum_{m=0}^{\infty} \sum_{m'=0}^m$ is adding all possible pairs (m, m') in the lower triangle, so we can add the terms vertically as the original sum indexing suggests, i.e.

$$\sum_{m=0}^{\infty} \sum_{m'=0}^m x_{(m,m')} = x_{(0,0)} + x_{(1,0)} + x_{(1,1)} + x_{(2,0)} + x_{(2,1)} + x_{(2,2)} + \dots, \quad (\text{A.121})$$

where the variable x is just a placeholder to indicate the order in which the sum is taking place. But we can also add the terms horizontally as

$$\sum_{m'=0}^{\infty} \sum_{m=m'}^{\infty} x_{(m,m')} = x_{(0,0)} + x_{(1,0)} + x_{(2,0)} + \dots + x_{(1,1)} + x_{(2,1)} + \dots, \quad (\text{A.122})$$

which still adds all of the lower triangle terms. Applying this reindexing results in

$$\lambda \sum_m z^m \sum_{m'=0}^m G_{m-m'} p(m', \tau) = \lambda \sum_{m'=0}^{\infty} \sum_{m=m'}^{\infty} z^m \theta (1 - \theta)^{m-m'} p(m', \tau), \quad (\text{A.123})$$

where we also substituted the definition of the geometric distribution $G_k = \theta(1-\theta)^k$. Redistributing the sums we can write

$$\lambda \sum_{m'=0}^{\infty} \sum_{m=m'}^{\infty} z^m \theta (1-\theta)^{m-m'} p(m', \tau) = \lambda \theta \sum_{n=0}^{\infty} (1-\theta)^{m'} P(m', \tau) \sum_{m=m'}^{\infty} [z(1-\theta)]^m. \quad (\text{A.124})$$

The next step requires us to look slightly ahead into what we expect to obtain. We are working on deriving an equation for the generating function $F(z, \tau)$ that when solved will allow us to compute what we care about, i.e. the probability function $p(m, \tau)$. Upon finding the function for $F(z, \tau)$, we will recover this probability distribution by evaluating derivatives of $F(z, \tau)$ at $z = 0$, whereas we can evaluate derivatives of $F(z, \tau)$ at $z = 1$ to instead recover the moments of the distribution. The point here is that when the dust settles we will evaluate z to be less than or equal to one. Furthermore, we know that the parameter of the geometric distribution θ must be strictly between zero and one. With these two facts we can safely state that $|z(1-\theta)| < 1$. Defining $n \equiv m - m'$ we rewrite the last sum in Eq. A.124 as

$$\sum_{m=m'}^{\infty} [z(1-\theta)]^m = \sum_{n=0}^{\infty} [z(1-\theta)]^{n+m'} \quad (\text{A.125})$$

$$= [z(1-\theta)]^{m'} \sum_{n=0}^{\infty} [z(1-\theta)]^n \quad (\text{A.126})$$

$$= [z(1-\theta)]^{m'} \left(\frac{1}{1 - z(1-\theta)} \right), \quad (\text{A.127})$$

where we use the geometric series since, as stated before, $|z(1-\theta)| < 1$. Putting these results together, the PDE for the generating function is

$$\frac{\partial F}{\partial \tau} = \frac{\partial F}{\partial z} - z \frac{\partial F}{\partial z} - \lambda F + \frac{\lambda \theta F}{1 - z(1-\theta)}. \quad (\text{A.128})$$

Changing variables to $\xi = 1 - \theta$ and simplifying gives

$$\frac{\partial F}{\partial \tau} + (z-1) \frac{\partial F}{\partial z} = \frac{(z-1)\xi}{1 - z\xi} \lambda F. \quad (\text{A.129})$$

Steady-state

To get at the mRNA distribution at steady state we first must solve Eq. A.129 setting the time derivative to zero. At steady-state, the PDE reduces to the ODE

$$\frac{dF}{dz} = \frac{\xi}{1 - z\xi} \lambda F, \quad (\text{A.130})$$

which we can integrate as

$$\int \frac{dF}{F} = \int \frac{\lambda \xi dz}{1 - \xi z}. \quad (\text{A.131})$$

The initial conditions for generating functions can be subtle and confusing. The key fact follows from the definition $F(z, t) = \sum_m z^m p(m, t)$. Clearly normalization of the distribution requires that $F(z = 1, t) = \sum_m p(m, t) = 1$. A subtlety is that sometimes the generating function may be undefined at $z = 1$, in which case the limit as z approaches 1 from below suffices to define the normalization condition. We also warn the reader that, while it is frequently convenient to change variables from z to a different independent variable, one must carefully track how the normalization condition transforms.

Continuing on, we evaluate the integrals (producing a constant c) which gives

$$\ln F = -\lambda \ln(1 - \xi z) + c \quad (\text{A.132})$$

$$F = \frac{c}{(1 - \xi z)^\lambda}. \quad (\text{A.133})$$

Only one choice for c can satisfy initial conditions, producing

$$F(z) = \left(\frac{1 - \xi}{1 - \xi z} \right)^\lambda = \left(\frac{\theta}{1 - z(1 - \theta)} \right)^\lambda, \quad (\text{A.134})$$

Recovering the steady-state probability distribution

To obtain the steady state mRNA distribution $p(m)$ we are aiming for we need to extract it from the generating function

$$F(z) = \sum_m z^m p(m). \quad (\text{A.135})$$

Taking a derivative with respect to z results in

$$\frac{dF(z)}{dz} = \sum_m m z^{m-1} p(m). \quad (\text{A.136})$$

Setting $z = 0$ leaves one term in the sum when $m = 1$

$$\left. \frac{dF(z)}{dz} \right|_{z=0} = \left(0 \cdot 0^{-1} \cdot p(0) + 1 \cdot 0^0 \cdot p(1) + 2 \cdot 0^1 \cdot p(2) + \dots \right) = p(1), \quad (\text{A.137})$$

since in the limit $\lim_{x \rightarrow 0^+} x^x = 1$. A second derivative of the generating function would result in

$$\frac{d^2 F(z)}{dz^2} = \sum_{m=0}^{\infty} m(m-1) z^{m-2} p(m). \quad (\text{A.138})$$

Again evaluating at $z = 0$ gives

$$\left. \frac{d^2 F(z)}{dz} \right|_{z=0} = 2p(z). \quad (\text{A.139})$$

In general any $p(m)$ is obtained from the generating function as

$$p(m) = \frac{1}{m!} \left. \frac{d^m F(z)}{dz} \right|_{z=0}. \quad (\text{A.140})$$

Let's now look at the general form of the derivative for our generating function in Eq. A.134. For $p(0)$ we simply evaluate $F(z = 0)$ directly, obtaining

$$p(0) = F(z = 0) = \theta^\lambda. \quad (\text{A.141})$$

The first derivative results in

$$\begin{aligned} \frac{dF(z)}{dz} &= \theta^\lambda \frac{d}{dz} (1 - z(1 - \theta))^{-\lambda} \\ &= \theta^\lambda [-\lambda(1 - z(1 - \theta))^{-\lambda-1} \cdot (\theta - 1)] \\ &= \theta^\lambda [\lambda(1 - z(1 - \theta))^{-\lambda-1} (1 - \theta)]. \end{aligned} \quad (\text{A.142})$$

Evaluating this at $z = 0$ as required to get $p(1)$ gives

$$\left. \frac{dF(z)}{dz} \right|_{z=0} = \theta^\lambda \lambda (1 - \theta) \quad (\text{A.143})$$

For the second derivative we find

$$\frac{d^2 F(z)}{dz^2} = \theta^\lambda [\lambda(\lambda + 1)(1 - z(1 - \theta))^{-\lambda-2} (1 - \theta)^2]. \quad (\text{A.144})$$

Again evaluating $z = 0$ gives

$$\left. \frac{d^2 F(z)}{dz^2} \right|_{z=0} = \theta^\lambda \lambda(\lambda + 1)(1 - \theta)^2. \quad (\text{A.145})$$

Let's go for one more derivative to see the pattern. The third derivative of the generating function gives

$$\frac{d^3 F(z)}{dz^3} = \theta^\lambda [\lambda(\lambda + 1)(\lambda + 2)(1 - z(1 - \theta))^{-\lambda-3} (1 - \theta)^3], \quad (\text{A.146})$$

which again we evaluate at $z = 0$

$$\left. \frac{d^3 F(z)}{dz^3} \right|_{z=0} = \theta^\lambda [\lambda(\lambda + 1)(\lambda + 2)(1 - \theta)^3]. \quad (\text{A.147})$$

If λ was an integer we could write this as

$$\left. \frac{d^3 F(z)}{dz^3} \right|_{z=0} = \frac{(\lambda+2)!}{(\lambda-1)!} \theta^\lambda (1-\theta)^3. \quad (\text{A.148})$$

Since λ might not be an integer we can write this using Gamma functions as

$$\left. \frac{d^3 F(z)}{dz^3} \right|_{z=0} = \frac{\Gamma(\lambda+3)}{\Gamma(\lambda)} \theta^\lambda (1-\theta)^3. \quad (\text{A.149})$$

Generalizing the pattern we then have that the m -th derivative takes the form

$$\left. \frac{d^m F(z)}{dz^m} \right|_{z=0} = \frac{\Gamma(\lambda+m)}{\Gamma(\lambda)} \theta^\lambda (1-\theta)^m. \quad (\text{A.150})$$

With this result we can use Eq. A.140 to obtain the desired steady-state probability distribution function

$$p(m) = \frac{\Gamma(m+\lambda)}{\Gamma(m+1)\Gamma(\lambda)} \theta^\lambda (1-\theta)^m. \quad (\text{A.151})$$

Note that the ratio of gamma functions is often expressed as a binomial coefficient, but since λ may be non-integer, this would be ill-defined. Re-expressing this exclusively in our variables of interest, burst rate λ and mean burst size b , we have

$$p(m) = \frac{\Gamma(m+\lambda)}{\Gamma(m+1)\Gamma(\lambda)} \left(\frac{1}{1+b} \right)^\lambda \left(\frac{b}{1+b} \right)^m. \quad (\text{A.152})$$

Adding repression

Deriving the generating function for mRNA distribution

Let us move from a one-state promoter to a two-state promoter, where one state has repressor bound and the other produces transcriptional bursts as above. A schematic of this model is shown as model 5 in Figure 2.1(C). Although now we have an equation for each promoter state, otherwise the master equation reads similarly to the one-state case, except with additional terms corresponding to transitions between promoter states, namely

$$\frac{d}{dt} p_R(m, t) = k_R^+ p_A(m, t) - k_R^- p_R(m, t) + (m+1) \gamma p_R(m+1, t) - m \gamma p_R(m, t) \quad (\text{A.153})$$

$$\begin{aligned} \frac{d}{dt} p_A(m, t) = & -k_R^+ p_A(m, t) + k_R^- p_R(m, t) + (m+1) \gamma p_A(m+1, t) - m \gamma p_A(m, t) \\ & - k_i p_A(m, t) + k_i \sum_{m'=0}^m \theta (1-\theta)^{m-m'} p_A(m', t), \end{aligned} \quad (\text{A.154})$$




Figure A.4: **Schematic of three-stage promoter from [10].** Adapted from Shahrezaei & Swain [10]. In their paper they derive a closed form solution for the protein distribution. Our two-state bursty promoter at the mRNA level can be mapped into their solution with some relabeling.

where $p_R(m, t)$ is the probability of the system having m mRNA copies and having repressor bound to the promoter at time t , and p_A is an analogous probability to find the promoter without repressor bound. k_{R+} and k_{R-} are, respectively, the rates at which repressors bind and unbind to and from the promoter, and γ is the mRNA degradation rate. k_i is the rate at which bursts initiate, and as before, the geometric distribution of burst sizes has mean $b = (1 - \theta)/\theta$.

Interestingly, it turns out that this problem maps exactly onto the three-stage promoter model considered by Shahrezaei and Swain in [10], with relabelings. Their approximate solution for protein distributions amounts to the same approximation we make here in regarding the duration of mRNA synthesis bursts as instantaneous, so their solution for protein distributions also solves our problem of mRNA distributions. Let us examine the analogy more closely. They consider a two-state promoter, as we do here, but they model mRNA as being produced one at a time and degraded, with rates v_0 and d_0 . Then they model translation as occurring with rate v_1 , and protein degradation with rate d_1 as shown in Figure A.4. Now consider the limit where $v_1, d_0 \rightarrow \infty$ with their ratio v_1/d_0 held constant. v_1/d_0 resembles the average burst size of translation from a single mRNA: these are the rates of two Poisson processes that compete over a transcript, which matches the story of geometrically distributed burst sizes. In other words, in our bursty promoter model we can think of the parameter θ as determining one competing process to end the burst and $(1 - \theta)$ as a process wanting to continue the burst. So after taking this limit, on timescales slow compared to v_1 and d_0 , it appears that transcription events fire at rate v_0 and produce a geometrically distributed burst of translation of mean size v_1/d_0 , which intuitively matches the story we have told above for mRNA with variables relabeled.

To verify this intuitively conjectured mapping between our problem and the solution in [10], we continue with a careful solution for the mRNA distribution using probability generating functions, following the ideas sketched in [10]. It is natural to nondimensionalize rates in the problem by γ , or equivalently, this amounts to measuring time in units of γ^{-1} . We are also only interested in steady state, so we set the time derivatives to zero, giving

$$0 = k_R^+ p_A(m) - k_R^- p_R(m) + (m+1)p_R(m+1) - mp_R(m) \quad (\text{A.155})$$

$$0 = -k_R^+ p_A(m) + k_R^- p_R(m) + (m+1)p_A(m+1) - mp_A(m) - k_i p_A(m) + k_i \sum_{m'=0}^m \theta(1-\theta)^{m-m'} p_A(m'), \quad (\text{A.156})$$

where for convenience we kept the same notation for all rates, but these are now expressed in units of mean mRNA lifetime γ^{-1} .

The probability generating function is defined as before in the constitutive case, except now we must introduce a generating function for each promoter state,

$$f_A(z) = \sum_{m=0}^{\infty} z^m p_A(m), \quad f_R(z) = \sum_{m=0}^{\infty} z^m p_R(m). \quad (\text{A.157})$$

Our real objective is the generating function $f(z)$ that generates the mRNA distribution $p(m)$, independent of what state the promoter is in. But since $p(m) = p_A(m) + p_R(m)$, it follows too that $f(z) = f_A(z) + f_R(z)$.

As before we multiply both equations by z^m and sum over all m . Each individual term transforms exactly as did an analogous term in the constitutive case, so the coupled ODEs for the generating functions read

$$0 = k_R^+ f_A(z) - k_R^- f_R(z) + \frac{\partial}{\partial z} f_R(z) - z \frac{\partial}{\partial z} f_R(z) \quad (\text{A.158})$$

$$0 = -k_R^+ f_A(z) + k_R^- f_R(z) + \frac{\partial}{\partial z} f_A(z) - z \frac{\partial}{\partial z} f_A(z) - k_i f_A(z) + k_i \frac{\theta}{1-z(1-\theta)} f_A(z), \quad (\text{A.159})$$

and after changing variables $\xi = 1 - \theta$ as before and rearranging, we have

$$0 = k_R^+ f_A(z) - k_R^- f_R(z) + (1-z) \frac{\partial}{\partial z} f_R(z) \quad (\text{A.160})$$

$$0 = -k_R^+ f_A(z) + k_R^- f_R(z) + (1-z) \frac{\partial}{\partial z} f_A(z) + k_i \frac{(z-1)\xi}{1-z\xi} f_A(z), \quad (\text{A.161})$$

We can transform this problem from two coupled first-order ODEs to a single second-order ODE by solving for f_A in the first and plugging into the second, giving

$$0 = (1-z) \frac{\partial f_R}{\partial z} + \frac{1-z}{k_R^+} \left(k_R^- \frac{\partial f_R}{\partial z} + \frac{\partial f_R}{\partial z} + (z-1) \frac{\partial^2 f_R}{\partial z^2} \right) + \frac{k_i}{k_R^+} \frac{(z-1)\xi}{1-z\xi} \left(k_R^- f_R + (z-1) \frac{\partial f_R}{\partial z} \right), \quad (\text{A.162})$$

where, to reduce notational clutter, we have dropped the explicit z dependence of f_A and f_R . Simplifying we have

$$0 = \frac{\partial^2 f_R}{\partial z^2} - \left(\frac{k_i \xi}{1-z\xi} + \frac{1+k_R^-+k_R^+}{1-z} \right) \frac{\partial f_R}{\partial z} + \frac{k_i k_R^- \xi}{(1-z\xi)(1-z)} f_R. \quad (\text{A.163})$$

This can be recognized as the hypergeometric differential equation, with singularities at $z = 1$, $z = \xi^{-1}$, and $z = \infty$. The latter can be verified by a change of variables from z to $x = 1/z$, being careful with the chain rule, and noting that $z = \infty$ is a singular point if and only if $x = 1/z = 0$ is a singular point.

The standard form of the hypergeometric differential equation has its singularities at 0, 1, and ∞ , so to take advantage of the standard form solutions to this ODE, we first need to transform variables to put it into a standard form. However, this is subtle. While any such transformation should work in principle, the solutions are expressed most simply in the neighborhood of $z = 0$, but the normalization condition that we need to enforce corresponds to $z = 1$. The easiest path, therefore, is to find a change of variables that maps 1 to 0, ∞ to ∞ , and ξ^{-1} to 1. This is most intuitively done in two steps.

First map the $z = 1$ singularity to 0 by the change of variables $v = z - 1$, giving

$$0 = \frac{\partial^2 f_R}{\partial v^2} + \left(\frac{k_i \xi}{(1+v)\xi - 1} + \frac{1+k_R^-+k_R^+}{v} \right) \frac{\partial f_R}{\partial v} + \frac{k_i k_R^- \xi}{((1+v)\xi - 1)v} f_R. \quad (\text{A.164})$$

Now two singularities are at $v = 0$ and $v = \infty$. The third is determined by $(1+v)\xi - 1 = 0$, or $v = \xi^{-1} - 1$. We want another variable change that maps this third singularity to 1 (without moving 0 or infinity). Changing variables again to $w = \frac{v}{\xi^{-1}-1} = \frac{\xi}{1-\xi} v$ fits the bill. In other words, the combined change of variables

$$w = \frac{\xi}{1-\xi} (z-1) \quad (\text{A.165})$$

maps $z = \{1, \xi^{-1}, \infty\}$ to $w = \{0, 1, \infty\}$ as desired. Plugging in, being mindful of the chain rule and noting $(1+v)\xi - 1 = (1-\xi)(w-1)$ gives

$$0 = \left(\frac{\xi}{1-\xi} \right)^2 \frac{\partial^2 f_R}{\partial w^2} + \left(\frac{\xi k_i}{(1-\xi)(w-1)} + \frac{\xi(1+k_R^-+k_R^+)}{(1-\xi)w} \right) \frac{\xi}{1-\xi} \frac{\partial f_R}{\partial w} + \frac{k_i k_R^- \xi^2}{(1-\xi)^2 w(w-1)} f_R. \quad (\text{A.166})$$

This is close to the standard form of the hypergeometric differential equation, and some cancellation and rearrangement gives

$$0 = w(w-1) \frac{\partial^2 f_R}{\partial w^2} + (k_i w + (1 + k_R^- + k_R^+)(w-1)) \frac{\partial f_R}{\partial w} + k_i k_R^- f_R. \quad (\text{A.167})$$

and a little more algebra produces

$$0 = w(1-w) \frac{\partial^2 f_R}{\partial w^2} + (1 + k_R^- + k_R^+ - (1 + k_i + k_R^- + k_R^+)w) \frac{\partial f_R}{\partial w} - k_i k_R^- f_R, \quad (\text{A.168})$$

which is the standard form. From this we can read off the solution in terms of hypergeometric functions ${}_2F_1$ from any standard source, e.g. [11], and identify the conventional parameters in terms of our model parameters. We want the general solution in the neighborhood of $w = 0$ ($z = 1$), which for a homogeneous linear second order ODE must be a sum of two linearly independent solutions. More precisely,

$$f_R(w) = C^{(1)} {}_2F_1(\alpha, \beta, \delta; w) + C^{(2)} w^{1-\delta} {}_2F_1(1 + \alpha - \delta, 1 + \beta - \delta, 2 - \delta; w) \quad (\text{A.169})$$

with parameters determined by

$$\begin{aligned} \alpha\beta &= k_i k_R^- \\ 1 + \alpha + \beta &= 1 + k_i + k_R^- + k_R^+ \\ \delta &= 1 + k_R^- + k_R^+ \end{aligned} \quad (\text{A.170})$$

and constants $C^{(1)}$ and $C^{(2)}$ to be set by boundary conditions. Solving for α and β , we find

$$\begin{aligned} \alpha &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ + \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right) \\ \beta &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ - \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right) \\ \delta &= 1 + k_R^- + k_R^+. \end{aligned} \quad (\text{A.171})$$

Note that α and β are interchangeable in the definition of ${}_2F_1$ and differ only in the sign preceeding the radical. Since the normalization condition requires that f_R be finite at $w = 0$, we can immediately set $C^{(2)} = 0$ to discard the second solution. This is because all the rate constants are strictly positive, so $\delta > 1$ and therefore $w^{1-\delta}$ blows up as $w \rightarrow 0$. Now that we have f_R , we would like to find the generating

function for the mRNA distribution, $f(z) = f_A(z) + f_R(z)$. We can recover f_A from our solution for f_R , namely

$$f_A(z) = \frac{1}{k_R^+} \left(k_R^- f_R(z) + (z-1) \frac{\partial f_R}{\partial z} \right) \quad (\text{A.172})$$

or

$$f_A(w) = \frac{1}{k_R^+} \left(k_R^- f_R(w) + w \frac{\partial f_R}{\partial w} \right), \quad (\text{A.173})$$

where in the second line we transformed our original relation between f_R and f_A to our new, more convenient, variable w . Plugging our solution for $f_R(w) = C^{(1)} {}_2F_1(\alpha, \beta, \delta; w)$ into f_A , we will require the differentiation rule for ${}_2F_1$, which tells us

$$\frac{\partial f_R}{\partial w} = C^{(1)} \frac{\alpha\beta}{\delta} {}_2F_1(\alpha+1, \beta+1, \delta+1; w), \quad (\text{A.174})$$

from which it follows that

$$f_A(w) = \frac{C^{(1)}}{k_R^+} \left(k_R^- {}_2F_1(\alpha, \beta, \delta; w) + w \frac{\alpha\beta}{\delta} {}_2F_1(\alpha+1, \beta+1, \delta+1; w) \right) \quad (\text{A.175})$$

and therefore

$$f(w) = C^{(1)} \left(1 + \frac{k_R^-}{k_R^+} \right) {}_2F_1(\alpha, \beta, \delta; w) + w \frac{C^{(1)}}{k_R^+} \frac{\alpha\beta}{\delta} {}_2F_1(\alpha+1, \beta+1, \delta+1; w). \quad (\text{A.176})$$

To proceed, we need one of the (many) useful identities known for hypergeometric functions, in particular

$$w \frac{\alpha\beta}{\delta} {}_2F_1(\alpha+1, \beta+1, \delta+1; w) = (\delta-1) ({}_2F_1(\alpha, \beta, \delta-1; w) - {}_2F_1(\alpha, \beta, \delta; w)). \quad (\text{A.177})$$

Substituting this for the second term in $f(w)$, we find

$$f(w) = \frac{C^{(1)}}{k_R^+} \left[(k_R^+ + k_R^-) {}_2F_1(\alpha, \beta, \delta; w) + (\delta-1) ({}_2F_1(\alpha, \beta, \delta-1; w) - {}_2F_1(\alpha, \beta, \delta; w)) \right], \quad (\text{A.178})$$

and since $\delta-1 = k_R^+ + k_R^-$, the first and third terms cancel, leaving only

$$f(w) = C^{(1)} \frac{k_R^+ + k_R^-}{k_R^+} {}_2F_1(\alpha, \beta, \delta-1; w). \quad (\text{A.179})$$

Now we enforce normalization, demanding $f(w = 0) = f(z = 1) = 1$. ${}_2F_1(\alpha, \beta, \delta - 1; 0) = 1$, so we must have $C^{(1)} = k_R^+ / (k_R^+ + k_R^-)$ and consequently

$$f(w) = {}_2F_1(\alpha, \beta, k_R^+ + k_R^-; w). \quad (\text{A.180})$$

Recalling that the mean burst size $b = (1 - \theta)/\theta = \xi/(1 - \xi)$ and $w = \frac{\xi}{1 - \xi}(z - 1) = b(z - 1)$, we can transform back to the original variable z to find the tidy result

$$f(z) = {}_2F_1(\alpha, \beta, k_R^+ + k_R^-; b(z - 1)), \quad (\text{A.181})$$

with α and β given above by

$$\begin{aligned} \alpha &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ + \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right) \\ \beta &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ - \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right). \end{aligned} \quad (\text{A.182})$$

Finally we are in sight of the original goal. We can generate the steady-state probability distribution of interest by differentiating the generating function,

$$p(m) = m! \left. \frac{\partial^m}{\partial z^m} f(z) \right|_{z=0}, \quad (\text{A.183})$$

which follows easily from its definition. Some contemplation reveals that repeated application of the derivative rule used above will produce products of the form $\alpha(\alpha + 1)(\alpha + 2) \cdots (\alpha + m - 1)$ in the expression for $p(m)$ and similarly for β and δ . These resemble ratios of factorials, but since α, β , and δ are not necessarily integer, we should express the ratios using gamma functions instead. More precisely, one finds

$$p(m) = \frac{\Gamma(\alpha + m)\Gamma(\beta + m)\Gamma(k_R^+ + k_R^-)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(k_R^+ + k_R^- + m)} \frac{b^m}{m!} {}_2F_1(\alpha + m, \beta + m, k_R^+ + k_R^- + m; -b) \quad (\text{A.184})$$

which is finally the probability distribution we sought to derive.

Numerical considerations and recursion formulas

Generalities

We would like to carry out Bayesian parameter inference on FISH data from [8], using Eq. (A.184) as our likelihood. This requires accurate (and preferably fast) numerical evaluation of the hypergeometric function ${}_2F_1$, which is a notoriously hard problem [12, 13], and our particular needs here present an especial challenge as we show below.

The hypergeometric function is defined by its Taylor series as

$${}_2F_1(a, b, c; z) = \sum_{l=0}^{\infty} \frac{\Gamma(a+l)\Gamma(b+l)\Gamma(c)}{\Gamma(a)\Gamma(b)\Gamma(c+l)} \frac{z^l}{l!} \quad (\text{A.185})$$

for $|z| < 1$, and by analytic continuation elsewhere. If $z \lesssim 1/2$ and α and β are not too large (absolute value below 20 or 30), then the series converges quickly and an accurate numerical representation is easily computed by truncating the series after a reasonable number of terms. Unfortunately, we need to evaluate ${}_2F_1$ over mRNA copy numbers fully out to the tail of the distribution, which can easily reach 50, possibly 100. From Eq. (A.184), this means evaluating ${}_2F_1$ repeatedly for values of a , b , and c spanning the full range from $O(1)$ to $O(10^2)$, even if α , β , and δ in Eq. (A.184) are small, with the situation even worse if they are not small. A naive numerical evaluation of the series definition will be prone to overflow and, if any of $a, b, c < 0$, then some successive terms in the series have alternating signs which can lead to catastrophic cancellations.

One solution is to evaluate ${}_2F_1$ using arbitrary precision arithmetic instead of floating point arithmetic, e.g., using the `mpmath` library in Python. This is accurate but incredibly slow computationally. To quantify how slow, we found that evaluating the likelihood defined by Eq. (A.184) ~ 50 times (for a typical dataset of interest from [8], with m values spanning 0 to ~ 50) using arbitrary precision arithmetic is 100-1000 fold slower than evaluating a negative binomial likelihood for the corresponding constitutive promoter dataset.

To claw back $\gtrsim 30$ fold of that slowdown, we can exploit one of the many catalogued symmetries involving ${}_2F_1$. The solution involves recursion relations originally explored by Gauss, and studied extensively in [12, 13]. They are sometimes known as contiguous relations and relate the values of any set of 3 hypergeometric functions whose arguments differ by integers. To rephrase this symbolically, consider a set of hypergeometric functions indexed by an integer n ,

$$f_n = {}_2F_1(a + \epsilon_i n, b + \epsilon_j n, c + \epsilon_k n; z), \quad (\text{A.186})$$

for a fixed choice of $\epsilon_i, \epsilon_j, \epsilon_k \in \{0, \pm 1\}$ (at least one of $\epsilon_i, \epsilon_j, \epsilon_k$ must be nonzero, else the set of f_n would contain only a single element). Then there exist known recurrence relations of the form

$$A_n f_{n-1} + B_n f_n + C_n f_{n+1} = 0, \quad (\text{A.187})$$

where A_n , B_n , and C_n are some functions of a , b , c , and z . In other words, for fixed $\epsilon_i, \epsilon_j, \epsilon_k, a, b$, and c , if we can merely evaluate ${}_2F_1$ twice, say for n' and $n' - 1$, then we can easily and rapidly generate values for arbitrary n .

This provides a convenient solution for our problem: we need repeated evaluations of ${}_2F_1(a + m, b + m, c + m; z)$ for fixed a, b , and c and many integer values of m . The idea is that we can use arbitrary precision arithmetic to evaluate ${}_2F_1$ for just two particular values of m and then generate ${}_2F_1$ for the other 50-100 values of m using the recurrence relation. In fact there are even more sophisticated ways of utilizing the recurrence relations that might have netted another factor of 2 speed-up, and possibly as much as a factor of 10, but the method described here had already reduced the computation time to an acceptable $O(1 \text{ min})$, so these more sophisticated approaches did not seem worth the time to pursue.

However, there are two further wrinkles. The first is that a naive application of the recurrence relation is numerically unstable. Roughly, this is because the three term recurrence relations, like second order ODEs, admit two linearly independent solutions. In a certain eigenbasis, one of these solutions dominates the other as $n \rightarrow \infty$, and as $n \rightarrow -\infty$, the dominance is reversed. If we fail to work in this eigenbasis, our solution of the recurrence relation will be a mixture of these solutions and rapidly accumulate numerical error. For our purposes, it suffices to know that the authors of [13] derived the numerically stable solutions (so-called *minimal solutions*) for several possible choices of $\epsilon_i, \epsilon_j, \epsilon_k$. Running the recurrence in the proper direction using a minimal solution is numerically robust and can be done entirely in floating point arithmetic, so that we only need to evaluate ${}_2F_1$ with arbitrary precision arithmetic to generate the seed values for the recursion.

The second wrinkle is a corollary to the first. The minimal solutions are only minimal for certain ranges of the argument z , and not all of the 26 possible recurrence relations have minimal solutions for all z . This can be solved by using one of the many transformation formulae for ${}_2F_1$ to convert to a different recurrence relation that has a minimal solution over the required domain of z , although this can require some trial and error to find the right transformation, the right recurrence relation, and the right minimal solution.

Particulars

Let us now demonstrate these generalities for our problem of interest. In order to evaluate the probability distribution of our model, Eq. (A.184), we need to evaluate

hypergeometric functions of the form ${}_2F_1(\alpha + m, \beta + m, \delta + m; -b)$ for values of m ranging from 0 to $O(100)$. The authors of [13] did not derive a recursion relation for precisely this case. We could follow their methods and do so ourselves, but it is much easier to convert to a case that they did consider. The strategy is to look through the minimal solutions tabulated in [13] and search for a transformation we could apply to ${}_2F_1(\alpha + m, \beta + m, \delta + m; -b)$ that would place the m 's (the variable being incremented by the recursion) in the same arguments of ${}_2F_1$ as the minimal solution. After some “guess and check,” we found that the transformation

$${}_2F_1(\alpha + m, \beta + m, \delta + m; -b) = (1 + b)^{-\alpha - m} {}_2F_1\left(\alpha + m, \delta - \beta, \delta + m; \frac{b}{1 + b}\right), \quad (\text{A.188})$$

produces a ${}_2F_1$ on the right hand side that closely resembles the minimal solutions $y_{3,m}$ and $y_{4,m}$ in Eq. 4.3 in [13]. Explicitly, these solutions are

$$y_{3,m} \propto {}_2F_1(-\alpha' + \delta' - m, -\beta' + \delta', 1 - \alpha' - \beta' + \delta' - m; 1 - z) \quad (\text{A.189})$$

$$y_{4,m} \propto {}_2F_1(\alpha' + m, \beta', 1 + \alpha' + \beta' - \delta' + m; 1 - z), \quad (\text{A.190})$$

where we have omitted prefactors which are unimportant for now. Which of these two we should use depends on what values z takes on. Equating $1 - z = b/(1 + b)$ gives $z = 1/(1 + b)$, and since b is strictly positive, z is bounded between 0 and 1. From Eq. 4.5 in [13], $y_{4,m}$ is the minimal solution for real z satisfying $0 < z < 2$, so this is the only minimal solution we need.

Now that we have our minimal solution, what recurrence relation does it satisfy? Confusingly, the recurrence relation of which $y_{4,m}$ is a solution increments different arguments of ${}_2F_1$ that does $y_{4,m}$: it increments the first only, rather than first and third. This recurrence relation can be looked up, e.g., Eq. 15.2.10 in [11], which is

$$(\delta' - (\alpha' + m))f_{m-1} + (2(\alpha' + m) - \delta' + (\beta' - \alpha')z)f_m + \alpha'(z - 1)f_{m+1} = 0. \quad (\text{A.191})$$

Now we must solve for the parameters appearing in the recurrence relation in terms of our parameters, namely by setting

$$\begin{aligned} \alpha' &= \alpha \\ \beta' &= \delta - \beta \\ 1 + \alpha' + \beta' - \delta' &= \delta \\ 1 - z &= \frac{b}{1 + b} \end{aligned} \quad (\text{A.192})$$

and solving to find

$$\begin{aligned}
 \alpha' &= \alpha \\
 \beta' &= \delta - \beta \\
 \delta' &= 1 + \alpha - \beta \\
 z &= \frac{1}{1+b}.
 \end{aligned}
 \tag{A.193}$$

Finally we have everything we need. The minimal solution

$$y_{4,m} = \frac{\Gamma(1 + \alpha' - \delta' + m)}{\Gamma(1 + \alpha' + \beta' - \delta' + m)} \times {}_2F_1(\alpha' + m, \beta', 1 + \alpha' + \beta' - \delta' + m; 1 - z),
 \tag{A.194}$$

where we have now included the necessary prefactors, is a numerically stable solution of the recurrence relation Eq. (A.191) if the recursion is run from large m to small m .

Let us finally outline the complete procedure as an algorithm to be implemented:

1. Compute the value of ${}_2F_1$ for the two largest m values of interest using arbitrary precision arithmetic.
2. Compute the prefactors to construct $y_{4,\max(m)}$ and $y_{4,\max(m)-1}$.
3. Recursively compute $y_{4,m}$ for all m less than $\max(m)$ down to $m = 0$.
4. Cancel off the prefactors of the resulting values of $y_{4,m}$ for all m to produce ${}_2F_1$ for all desired m values.

With ${}_2F_1$ computed, the only remaining numerical danger in computing $p(m)$ in Eq. (A.184) is overflow of the gamma functions. This is easily solved by taking the log of the entire expression and using standard routines to compute the log of the gamma functions, then exponentiating the entire expression at the end if $p(m)$ is needed rather than $\log p(m)$.

A.3 Bayesian inference

The problem of parameter inference

One could argue that the whole goal of formulating theoretical models about nature is to sharpen our understanding from qualitative statements to precise quantitative assertions about the relevant features of the natural phenomena in question [14].

It is in these models that we intend to distill the essential parts of the object of study. Writing down such models leads to a propagation of mathematical variables that parametrize our models. By assigning numerical values to these parameters we can compute concrete predictions that can be contrasted with experimental data. For these predictions to match the data the parameter values have to carefully be chosen from the whole parameter space. But how do we go about assessing the effectiveness of different regions of parameter space to speak to the ability of our model to reproduce the experimental observations? The language of probability, and more specifically of Bayesian statistics is – we think – the natural language to tackle this question.

Bayes' theorem

Bayes' theorem is a simple mathematical statement that can apply to *any* logical conjecture. For two particular events A and B that potentially depend on each other Bayes' theorem gives us a recipe for how to update our beliefs about one, let us say B , given some state of knowledge, or lack thereof, about A . In its most classic form Bayes' theorem is written as

$$P(B | A) = \frac{P(A | B)P(B)}{P(A)}, \quad (\text{A.195})$$

where the vertical line $|$ is read as “given that”. So $P(B | A)$ is read as probability of B given that A took place. A and B can be any logical assertion. In particular the problem of Bayesian inference focuses on the question of finding the probability distribution of a particular parameter value given the data.

For a given model with a set of parameters $\vec{\theta} = (\theta_1, \theta_2, \dots, \theta_n)$, the so-called *posterior distribution* $P(\vec{\theta} | D)$, where D is the experimental data, quantifies the plausibility of a set of parameter values given our observation of some particular dataset. In other words, through the application of Bayes' formula we update our beliefs on the possible values that parameters can take upon learning the outcome of a particular experiment. We specify the word “update” as we come to every inference problem with prior information about the plausibility of particular regions of parameter space even before performing any experiment. Even when we claim as researchers that we are totally ignorant about the values that the parameters in our models can take, we always come to a problem with domain expertise that can be exploited. If this was not the case, it is likely that the formulation of our model is not going to capture the phenomena we claim to want to understand. This prior

information is captured in the *prior probability* $P(\vec{\theta})$. The relationship between how parameter values can connect with the data is encoded in the *likelihood function* $P(D | \vec{\theta})$. Our theoretical model, whether deterministic or probabilistic, is encoded in this term that can be intuitively understood as the probability of having observed the particular experimental data we have at hand given that our model is parametrized with the concrete values $\vec{\theta}$. Implicitly here we are also conditioning on the fact that our theoretical model is “true,” i.e. the model itself if evaluated or simulated in the computer is capable of generating equivalent datasets to the one we got to observe in an experiment. In this way Bayesian inference consists of applying Bayes’ formula as

$$P(\vec{\theta} | D) \propto P(D | \vec{\theta})P(\vec{\theta}). \quad (\text{A.196})$$

Notice that rather than writing the full form of Bayes’ theorem, we limit ourselves to the terms that depend on our quantity of interest – that is the parameter values themselves $\vec{\theta}$ – as the denominator $P(D)$ only serves as a normalization constant.

We also emphasize that the dichotomy we have presented between prior and likelihood is more subtle. Although it is often stated that our prior knowledge is entirely encapsulated by the obviously named prior probability $P(\vec{\theta})$, this is usually too simplistic. The form(s) we choose for our likelihood function $P(D | \vec{\theta})$ also draw heavily on our prior domain expertise and the assumptions, implicit and explicit, that these choices encode are at least as important, and often inseparable from, the prior probability, as persuasively argued in [15].

The likelihood function

As we alluded in the previous section it is through the likelihood function $P(D | \vec{\theta})$ that we encode the connection between our parameter values and the experimental observables. Broadly speaking there are two classes of models that we might need to encode into our likelihood function:

- **Deterministic models:** Models for which a concrete selection of parameter values give a single output. Said differently, models with a one-to-one mapping between inputs and outputs.
- **Probabilistic models:** As the name suggests, models that, rather than having a one-to-one input-output mapping, describe the full probability distribution of possible outputs.

In this paper we focus on inference done with probabilistic models. After all, the chemical master equations we wrote down describe the time evolutions of the mRNA probability distribution. So all our terms $P(\vec{\theta} \mid D)$ will be given by the steady-state solution of the corresponding chemical master equation in question. This is rather convenient as we do not have to worry about adding a statistical model on top of our model to describe deviations from the predictions. Instead our models themselves focus on predicting such variation in cell count.

Prior selection

The different models explored in this work embraced different levels of coarse-graining that resulted in a diverse number of parameters for different models. For each of these model configurations Bayes' theorem demands from us to represent our preconceptions on the possible parameter values in the form of the prior $P(\vec{\theta})$. Throughout this work for models with > 1 parameter we assign independent priors to each of the parameters; this is

$$P(\vec{\theta}) = \prod_{i=1}^n P(\theta_i). \quad (\text{A.197})$$

Although it is not uncommon practice to use non-informative, or maximally uninformative priors, we are of the mindset that this is a disservice to the philosophical and practical implications of Bayes' theorem. It sounds almost contradictory to claim that can we represent our thinking about a natural phenomenon in the form of a mathematical model – in the context of Bayesian inference this means choosing a form for the likelihoods, and even making this choice presupposes prior understanding or assumptions as to the relevant features in the system under study – but that we have absolutely no idea what the parameter values could or could not be. We therefore make use of our own expertise, many times in the form of order-of-magnitude estimates, to write down weakly-informative prior distributions for our parameters.

For our particular case all of the datasets from [8] used in this paper have $O(10^3)$ data points. What this implies is that our particular choice of priors will not significantly affect our inference as long as they are broad enough. A way to see why this is the case is to simply look at Bayes' theorem. For $N = 1000 - 3000$ datum all of the independent of each other and $n \ll 10^3$ parameters Bayes' theorem reads as

$$P(\vec{\theta} \mid D) \propto \prod_{k=1}^N P(d_k \mid \vec{\theta}) \prod_{i=1}^n P(\theta_i), \quad (\text{A.198})$$

where d_k represents the k -th datum. That means that if our priors span a wide range of parameter space, the posterior distribution would be dominated by the likelihood function.

Expectations and marginalizations

For models with more than one or two parameters, it is generally difficult to visualize or reason about the full joint posterior distribution $P(\vec{\theta} \mid D)$ directly. One of the great powers of Bayesian analysis is *marginalization*, allowing us to reduce the dimensionality to only the parameters of immediate interest by averaging over the other dimensions. Formally, for a three dimensional model with parameters θ_1 , θ_2 , and θ_3 , we can for instance marginalize away θ_3 to produce a 2D posterior as

$$P(\theta_1, \theta_2 \mid D) \propto \int_{\theta_3} d\theta_3 P(\theta_1, \theta_2, \theta_3 \mid D), \quad (\text{A.199})$$

or we can marginalize away θ_1 and θ_3 to produce the 1D marginal posterior of θ_2 alone, which would be

$$P(\theta_2 \mid D) \propto \int_{\theta_1} d\theta_1 \int_{\theta_3} d\theta_3 P(\theta_1, \theta_2, \theta_3 \mid D). \quad (\text{A.200})$$

Conceptually, this is what we did in generating the 2D slices of the full 9D model in Figure 2.4(A). In practice, this marginalization is even easier with Markov Chain Monte Carlo samples in hand. Since each point is simply a list of parameter values, we simply ignore the parameters which we want to marginalize away [16].

Markov Chain Monte Carlo

The theory and practice of Bayesian inference with Markov Chain Monte Carlo (MCMC) is a rich subject with fascinating and deep analogies to statistical mechanics, even drawing on classical Hamiltonian mechanics and general relativity in its modern incarnations. We refer the interested reader to [16] and [17] for excellent introductions. Here we merely give a brief summary of the MCMC computations carried out in this work.

We used the Python package `emcee` for most of the MCMC sampling in this work. For the constitutive promoter inference, we also ran sampling with the excellent Stan modeling language as a check. We did not use Stan for the inference of the simple repression model because implementing the gradients of the hypergeometric function ${}_2F_1$ appearing in Eq. A.184, the probability distribution for our bursty

model with repression, would have been an immensely challenging task. `emcee` was more than adequate for our purposes, and we were perhaps lucky that the 9-D posterior model for the model of simple repression with bursty promoter was quite well behaved and did not require the extra power of the Hamiltonian Monte Carlo algorithm provided by Stan [18]. Source code for all statistical inference will be made available at https://github.com/RPGroup-PBoC/bursty_transcription.

Bayesian inference on constitutive promoters

Having introduced the ideas behind Bayesian inference we are ready to apply the theoretical machinery to our non-equilibrium models. In particular in this section we will focus on model 1 and model 5 in Figure 2.2(A). Model 1, the Poisson promoter, will help us build practical intuition into the implementation of the Bayesian inference pipeline as we noted in Section 2.4 of the main text that this model cannot be reconciled with experimental data from observables such as the Fano factor. In other words, we acknowledge that this model is “wrong,” but we still see value in going through the analysis since the simple nature of the model translates into a neat statistical analysis.

Model 1 - Poisson promoter

As specified in the main text, the mRNA steady-state distribution for model 1 in Figure 2.2(A) is Poisson with parameter λ . Throughout this Appendix we will appeal to the convenient notation for probability distributions of the form

$$m \sim \text{Poisson}(\lambda), \quad (\text{A.201})$$

where the symbol “ \sim ” can be read as *is distributed according to*. So the previous equation can be read as: the mRNA copy number m is distributed according to a Poisson distribution with parameter λ . Our objective then is to compute the posterior probability distribution $P(\lambda \mid D)$, where, as in the main text, $D = \{m_1, m_2, \dots, m_N\}$ are the data consisting of single-cell mRNA counts. Since we can assume that each of the cells mRNA counts are independent of any other cells, our likelihood function $P(D \mid \lambda)$ consists of the product of N Poisson distributions.

To proceed with the inference problem we need to specify a prior. In this case we are extremely data-rich, as the dataset from Jones et. al [8] has of order 1000-3000 single-cell measurements for each promoter, so our choice of prior matters little here, as long as it is sufficiently broad. A convenient choice for our problem is to use a *conjugate* prior. A conjugate prior is a special prior that causes the posterior to have

the same functional form as the prior, simply with updated model parameters. This makes calculations analytically tractable and also offers a nice interpretation of the inference procedure as updating our knowledge about the model parameters. This makes conjugate priors very useful when they exist. The caveat is that conjugate priors only exist for a very limited number of likelihoods, mostly with only one or two model parameters, so in almost all other Bayesian inference problems, we must tackle the posterior numerically.

But, for the problem at hand, a conjugate prior does in fact exist. For a Poisson likelihood of identical and identically distributed data, the conjugate prior is a gamma distribution, as can be looked up in, e.g., [16], Section 2.6. Putting a gamma prior on λ introduces two new parameters α and β which parametrize the gamma distribution itself, which we use to encode the range of λ values we view as reasonable. Recall λ is the mean steady-state mRNA count per cell, which *a priori* could plausibly be anywhere from 0 to a few hundred. $\alpha = 1$ and $\beta = 1/50$ achieve this, since the gamma distribution is strictly positive with mean α/β and standard deviation $\sqrt{\alpha}/\beta$. To be explicit, then, our prior is

$$\lambda \sim \text{Gamma}(\alpha, \beta) \quad (\text{A.202})$$

As an aside, note that if we did not know that our prior was a conjugate prior, we could still write down our posterior distribution from its definition as

$$p(\lambda \mid D, \alpha, \beta) \propto p(D \mid \lambda) p(\lambda \mid \alpha, \beta) \propto \left(\prod_{k=1}^N \frac{\lambda^{m_k} e^{-\lambda}}{m_k!} \right) \frac{\beta}{\Gamma(\alpha)} (\beta\lambda)^{\alpha-1} e^{-\beta\lambda}. \quad (\text{A.203})$$

Without foreknowledge that this in fact reduces to a gamma distribution, this expression might appear rather inscrutable. When conjugate priors are unavailable for the likelihood of interest - which is almost always the case for models with > 1 model parameter - this inscrutability is the norm, and making sense of posteriors analytically is almost always impossible. Fortunately, MCMC sampling provides us a powerful method of constructing posteriors numerically which we will make use of extensively.

Since we did use a conjugate prior, we may simply look up our posterior in any standard reference such as [16], Section 2.6, from which we find that

$$\lambda \sim \text{Gamma}(\alpha + \bar{m}N, \beta + N), \quad (\text{A.204})$$

where we defined the sample mean $\bar{m} = \frac{1}{N} \sum_k m_k$ for notational convenience. A glance at the FISH data from [8] reveals that N is $O(10^3)$ and $\langle m \rangle \gtrsim 0.1$ for all

constitutive strains in [8], so $\bar{m}N \gtrsim 10^2$. Therefore as we suspected, our prior parameters are completely overwhelmed by the data. The prior behaves, in a sense, like β extra “data points” with a mean value of $(\alpha - 1)/\beta$ [16], which gives us some intuition for how much data is needed to overwhelm the prior in this case: enough data N such that $\beta \ll N$ and $\alpha/\beta \ll \bar{m}$. In fact, $\bar{m}N$ and N are so large that we can, to an excellent approximation, ignore the α and β dependence and approximate the gamma distribution as a Gaussian with mean \bar{m} and standard deviation $\sqrt{\bar{m}/N}$, giving

$$\lambda \sim \text{Gamma}(\alpha + \bar{m}N, \beta + N) \approx \text{Normal}\left(\bar{m}, \sqrt{\frac{\bar{m}}{N}}\right). \quad (\text{A.205})$$

As an example with real numbers, for the *lacUV5* promoter, Jones et. al [8] measured 2648 cells with an average mRNA count per cell of $\bar{m} \approx 18.7$. In this case then, our posterior is

$$\lambda \sim \text{Normal}(18.7, 0.08), \quad (\text{A.206})$$

which suggests we have inferred our model’s one parameter to a precision of order 1%.

This is not wrong, but it is not the full story. The model’s posterior distribution is tightly constrained, but is it a good generative model? In other words, if we use the model to generate synthetic data in the computer does it generate data that look similar to our actual data, and is it therefore plausible that the model captures the important features of the data generating process? This intuitive notion can be codified with *posterior predictive checks*, or PPCs, and we will see that this simple Poisson model fails badly.

The intuitive idea of posterior predictive checks is simple:

1. Make a random draw of the model parameter λ from the posterior distribution.
2. Plug that draw into the likelihood and generate a synthetic dataset $\{m_k\}$ conditioned on λ .
3. Repeat many times.

More formally, the posterior predictive distribution can be thought of as the distribution of future yet-to-be-observed data, conditioned on the data we have already observed. Clearly if those data appear quite different, the model has a problem. Put another way, if we suppose the generative model is true, i.e. we claim that our model

explains the process through which our observed experimental data was generated, then the synthetic datasets we generate should resemble the actual observed data. If this is not the case, it suggests the model is missing important features. All the data we consider in this work are 1D (distributions of mRNA counts over a population) so empirical cumulative distribution functions ECDFs are an excellent visual means of comparing synthetic and observed datasets. In general for higher dimensional datasets, much of the challenge is in merely designing good visualizations that can actually show if synthetic and observed data are similar or not.

For our example Poisson promoter model then, we merely draw many random numbers, say 1000, from the Gaussian posterior in Eq. A.206. For each one of those draws, we generate a dataset from the likelihood, i.e., we draw 2648 (the number of observed cells in the actual dataset) Poisson-distributed numbers for each of the 1000 posterior draws, for a total of 2648000 samples from the posterior predictive distribution.

To compare so many samples with the actual observed data, one excellent visualization for 1D data is ECDFs of the quantiles, as shown for our Poisson model in Figure 2.3(B) in the main text.

Model 5 - Bursty promoter

Let us now consider the problem of parameter inference from FISH data for model five from Figure 2.1(C). As derived in Appendix A.2, the steady-state mRNA distribution in this model is a negative binomial distribution, given by

$$p(m) = \frac{\Gamma(m + k_i)}{\Gamma(m + 1)\Gamma(k_i)} \left(\frac{1}{1 + b} \right)^{k_i} \left(\frac{b}{1 + b} \right)^m, \quad (\text{A.207})$$

where b is the mean burst size and k_i is the burst rate nondimensionalized by the mRNA degradation rate γ . As sketched earlier, we can intuitively think about this distribution through a simple story. The story of this distribution is that the promoter undergoes geometrically-distributed bursts of mRNA, where the arrival of bursts is a Poisson process with rate k_i and the mean size of a burst is b .

As for the Poisson promoter model, this expression for the steady-state mRNA distribution is exactly the likelihood we want to use in Bayes' theorem. Again denoting the single-cell mRNA count data as $D = \{m_1, m_2, \dots, m_N\}$, here Bayes' theorem takes the form

$$p(k_i, b \mid D) \propto p(D \mid k_i, b)p(k_i, b), \quad (\text{A.208})$$

where the likelihood $p(D \mid k_i, b)$ is given by the product of N negative binomials as in Eq. A.207. We only need to choose priors on k_i and b . For the datasets from [8] that we are analyzing, as for the Poisson promoter model above we are still data-rich so the prior's influence remains weak, but not nearly as weak because the dimensionality of our model has increased from one to two.

We follow the guidance of [16], Section 2.9 in opting for weakly-informative priors on k_i and b (conjugate priors do not exist for this problem), and we find “street-fighting estimates” [19] to be an ideal way of constructing such priors. The idea of weakly informative priors is to allow all remotely plausible values of model parameters while excluding the completely absurd or unphysical.

Consider k_i . Some of the strongest known bacterial promoters control rRNA genes and initiate transcripts no faster than $\sim 1/\text{sec}$. It would be exceedingly strange if any of the constitutive promoters from [8] were stronger than that, so we can take that as an upper bound. For a lower bound, if transcripts are produced too rarely, there would be nothing to see with FISH. The datasets for each strain contain of order 10^3 cells, and if the $\langle m \rangle = k_i b / \gamma \lesssim 10^{-2}$, then the total number of expected mRNA detections would be single-digits or less and we would have essentially no data on which to carry out inference. So assuming b is not too different from 1, justified next, and an mRNA lifetime of $\gamma^{-1} \sim 3 - 5$ min, this gives us soft bounds on k_i/γ of perhaps 10^{-2} and 3×10^1 .

Next consider mean burst size b . This parametrization of the geometric distribution allows bursts of size zero (which could represent aborted transcripts and initiations), but it would be quite strange for the mean burst size b to be below $\sim 10^{-1}$, for which nearly all bursts would be of size zero or one. For an upper bound, if transcripts are initiating at a rate somewhat slower than rRNA promoters, then it would probably take a time comparable to the lifetime of an mRNA to produce a burst larger than 10-20 transcripts, which would invalidate the approximation of the model that the duration of bursts are instantaneous compared to other timescales in the problem. So we will take soft bounds of 10^{-1} and 10^1 for b .

Note that the natural scale for these “street-fighting estimates” was a log scale. This is commonly the case that our prior sense of reasonable and unreasonable parameters is set on a log scale. A natural way to enforce these soft bounds is therefore to use a lognormal prior distribution, with the soft bounds set ± 2 standard deviations from the mean.

With this, we are ready to write our full generative model as

$$\begin{aligned}\ln k_i &\sim \text{Normal}(-0.5, 2), \\ \ln b &\sim \text{Normal}(0.5, 1), \\ m &\sim \text{NBinom}(k_i, b).\end{aligned}\tag{A.209}$$

Section 2.5 in the main text details the results of applying this inference to the single-cell mRNA counts data. There we show the posterior distribution for the two parameters for different promoters. Figure A.5 shows the so-called posterior predictive checks (see main text for explanation) for all 18 unregulated promoters shown in the main text.

Bayesian inference on the simple-repression architecture

As detailed in 2.5 in the main text the inference on the unregulated promoter served as a stepping stone towards our ultimate goal of inferring repressor rates from the steady-state mRNA distributions of simple-repression architectures. For this we expand the one-state bursty promoter model to a two-state promoter as schematized in Figure 2.1(C) as model 5. This model adds two new parameters: the repressor binding rate k^+ , solely function of the repressor concentration, and the repressor dissociation rate k^- , solely a function of the repressor-DNA binding affinity.

The structure of the data in [8] for regulated promoters tuned these two parameters independently. In their work the production of the LacI repressor was under the control of an inducible promoter regulated by the TetR repressor as schematized in Figure A.6. When TetR binds to the small molecule anhydrotetracycline (aTc), it shifts to an inactive conformation unable to bind to the DNA. This translates into an increase in gene expression level. In other words, the higher the concentration of aTc added to the media, the less TetR repressors that can control the expression of the *lacI* gene, so the higher the concentration of LacI repressors in the cell. So by tuning the amount of aTc in the media where the experimental strains were grown they effectively tune k^+ in our simple theoretical model. On the other hand to tune k^- the authors swap three different binding sites for the LacI repressor, each with different repressor-DNA binding affinities previously characterized [21].

What this means is that we have access to data with different combinations of k^- and k^+ . We could naively try to fit the kinetic parameters individually for each of the datasets, but there is no reason to believe that the binding site identity for the LacI repressor somehow affects its expression level controlled from a completely different location in the genome, nor vice versa. In other words, what makes the

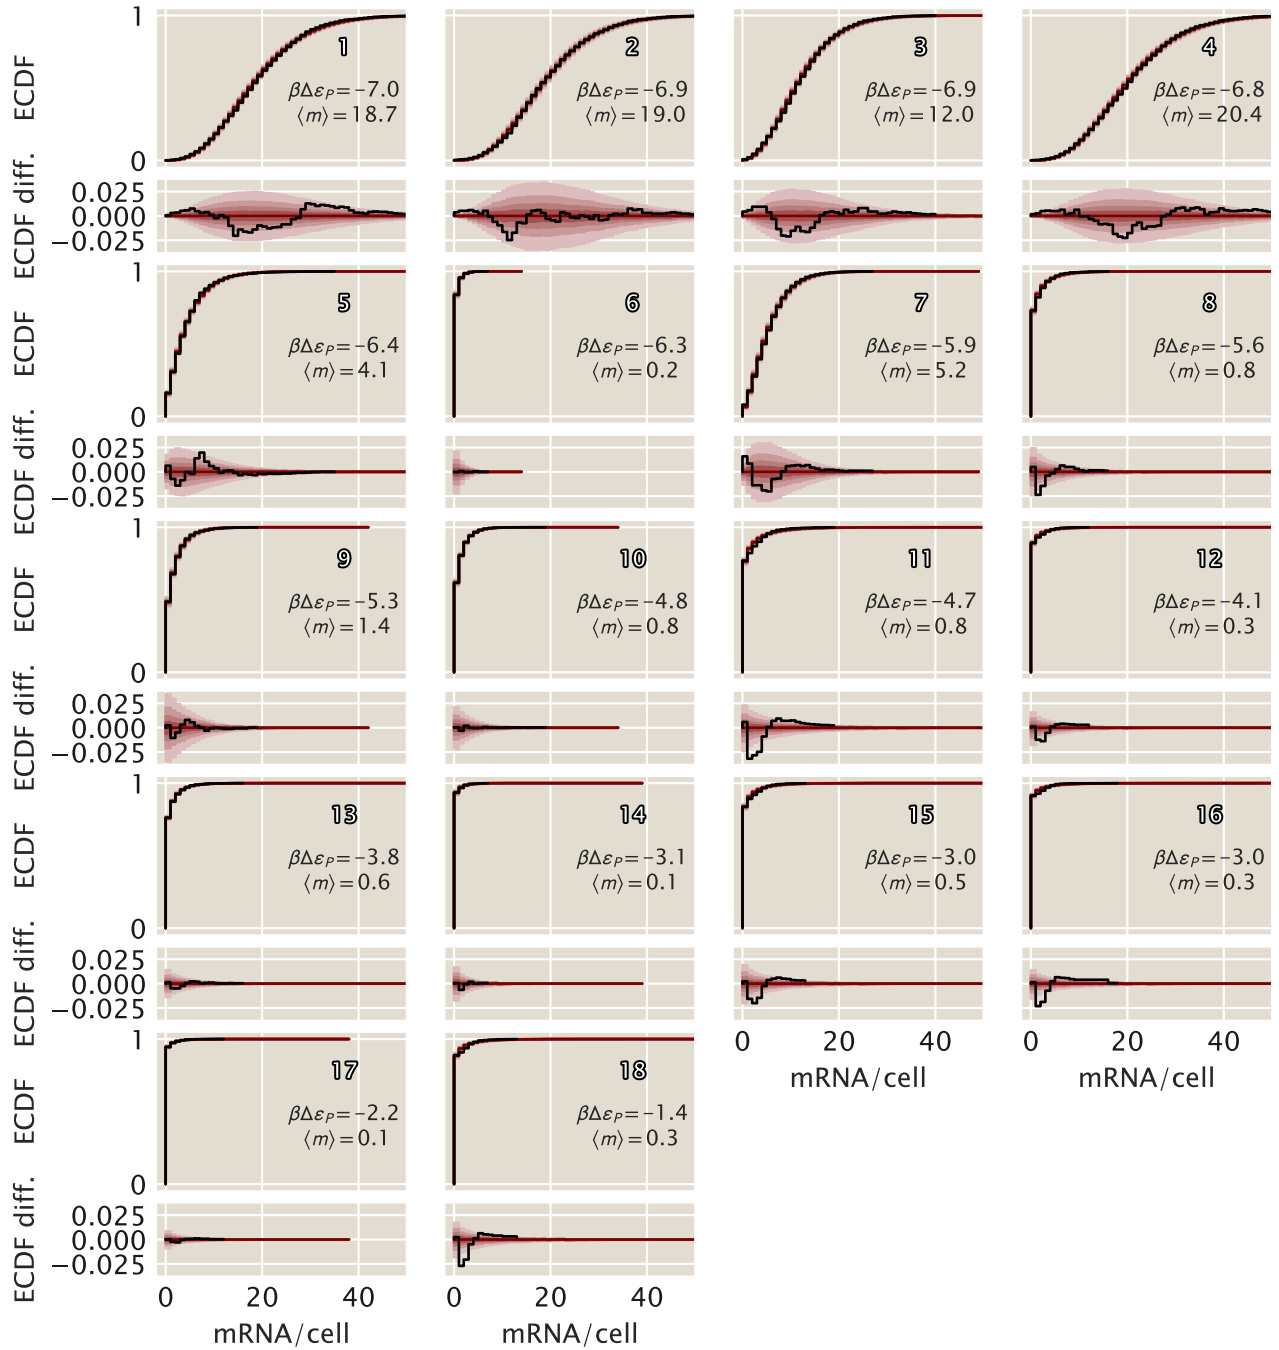


Figure A.5: Theory-data comparison of inference on unregulated promoters. Comparison of the inference (red shaded area) vs the experimental measurements (black lines) for 18 different unregulated promoters with different mean mRNA expression levels from Ref. [8]. Upper panels show the empirical cumulative distribution function (ECDF), while the lower panels show the differences with respect to the median of the posterior samples. White numbers are the same as in Figure 2.1 for cross comparison. The predicted binding energies $\beta\Delta\epsilon_p$ were obtained from the energy matrix model in Ref. [20]




Figure A.6: **aTc controlled expression of LacI repressor.** Schematic of the circuit used in [8] to control the expression of the LacI repressor. The *lacI* gene is under the control of the TetR repressor. As the TetR repressor is inactivated upon binding of anhydrotetracycline or aTc, the more aTc added to the media were cells are growing, the less TetR repressors available to control the expression of the *lacI* gene, resulting in more LacI repressors per cell. LacI simultaneously controls the expression of the mRNA on which single-molecule mRNA FISH was performed for gene expression quantification.

most sense it to fit all datasets together to obtain a single value for each of the association and dissociation rates. What this means, as described in Section 2.5 of the main text is that we have a seven dimensional parameter space with four possible association rates k^+ given the four available aTc concentrations, and three possible dissociation rates k^- given the three different binding sites available in the dataset.

Formally now, denote the set of seven repressor rates to be inferred as

$$\vec{k} = \{k_{Oid}^-, k_{O1}^-, k_{O2}^-, k_{0.5}^+, k_1^+, k_2^+, k_{10}^+\}. \quad (\text{A.210})$$

Note that since the repressor copy numbers are not known directly as explained before, we label their association rates by the concentration of aTc. Bayes theorem reads simply

$$p(\vec{k}, k_i, b \mid D) \propto p(D \mid \vec{k}, k_i, b) p(\vec{k}, k_i, b), \quad (\text{A.211})$$

where D is the set of all N observed single-cell mRNA counts across the various conditions. We assume that individual single-cell measurements are independent so that the likelihood factorizes as

$$p(D \mid \vec{k}, k_i, b) = \prod_{j=1}^N p(m \mid \vec{k}, k_i, b) = \prod_{j=1}^N p(m \mid k_j^+, k_j^-, k_i, b) \quad (\text{A.212})$$

where k_j^\pm represent the appropriate binding and unbinding rates for the j -th measured cell. Our likelihood function, previously derived in Appendix A.2, is given by the rather complicated result in Eq. A.184, which for completeness we reproduce here

as

$$p(m \mid k_R^+, k_R^-, k_i, b) = \frac{\Gamma(\alpha + m)\Gamma(\beta + m)\Gamma(k_R^+ + k_R^-)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(k_R^+ + k_R^- + m)} \frac{b^m}{m!} \times {}_2F_1(\alpha + m, \beta + m, k_R^+ + k_R^- + m; -b). \quad (\text{A.213})$$

where α and β , defined for notational convenience, are

$$\begin{aligned} \alpha &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ + \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right) \\ \beta &= \frac{1}{2} \left(k_i + k_R^- + k_R^+ - \sqrt{(k_i + k_R^- + k_R^+)^2 - 4k_i k_R^-} \right). \end{aligned} \quad (\text{A.214})$$

Next we specify priors. As for the constitutive model, weakly informative lognormal priors are a natural choice for all our rates. We found that if the priors were too weak, our MCMC sampler would often become stuck in regions of parameter space with very low probability density, unable to move. We struck a balance in choosing our prior widths between helping the sampler run while simultaneously verifying that the marginal posteriors for each parameter were not artificially constrained or distorted by the presence of the prior. The only exception to this is the highly informative priors we placed on k_i and b , since we have strong knowledge of them from our inference of constitutive promoters above.

With priors and likelihood specified we may write down our complete generative model as

$$\begin{aligned} \log_{10} k_i &\sim \text{Normal}(0.725, 0.025) \\ \log_{10} b &\sim \text{Normal}(0.55, 0.025) \\ \log_{10} k_{0.5}^+ &\sim \text{Normal}(-0.45, 0.3) \\ \log_{10} k_1^+ &\sim \text{Normal}(0.6, 0.3) \\ \log_{10} k_2^+ &\sim \text{Normal}(1.15, 0.3) \\ \log_{10} k_{10}^+ &\sim \text{Normal}(1.5, 0.3) \\ \log_{10} k_{Oid}^- &\sim \text{Normal}(-0.25, 0.3) \\ \log_{10} k_{O1}^- &\sim \text{Normal}(0.1, 0.3) \\ \log_{10} k_{O2}^- &\sim \text{Normal}(0.45, 0.3) \\ m &\sim \text{Likelihood}(k_R^+, k_R^-, k_i, b), \end{aligned} \quad (\text{A.215})$$

where the likelihood is specified by Eq. A.213. We ran MCMC sampling on the full nine dimensional posterior specified by this generative model.

We found that fitting a single operator/aTc concentration at a time with a single binding and unbinding rate did not yield a stable inference for most of the possible operator/aTc combinations. In other words, a single dataset could not independently resolve the binding and unbinding rates, only their ratio as set by the mean fold-change in Figure 2.1 in the main text. Only by making the assumption of a single unique binding rate for each repressor copy number and a single unique unbinding rate for each binding site, as done in Figure 2.4(A), was it possible to independently resolve the rates and not merely their ratios.

We also note that we found it necessary to exclude the very weakly and very strongly repressed datasets from Jones et. al. [8]. In both cases there was, in a sense, not enough information in the distributions for our inference algorithm to extract, and their inclusion simply caused problems for the MCMC sampler without yielding any new insight. For the strongly repressed data (Oid, 10 ng/mL aTc), with > 95% of cells with zero mRNA, there was quite literally very little data from which to infer rates. And the weakly repressed data, all with the repressor binding site O3, had an unbinding rate so fast that the sampler essentially sampled from the prior; the likelihood had negligible influence, meaning the data was not informing the sampler in any meaningful way, so no inference was possible.

References

- ¹A. Sanchez, H. G. Garcia, D. Jones, R. Phillips, and J. Kondev, “Effect of promoter architecture on the cell-to-cell variability in gene expression”, *PLoS computational biology* **7**, e1001100 (2011).
- ²M. Razo-Mejia, S. Marzen, G. Chure, R. Taubman, M. Morrison, and R. Phillips, “First-principles prediction of the information processing capacity of a simple genetic circuit”, arXiv:2005.03214 [q-bio] (2020).
- ³E. L. King and C. Altman, “A Schematic Method of Deriving the Rate Laws for Enzyme-Catalyzed Reactions”, *The Journal of Physical Chemistry* **60**, 1375–1378 (1956).
- ⁴T. L. Hill, “Studies in irreversible thermodynamics IV. diagrammatic representation of steady state fluxes for unimolecular systems”, *Journal of Theoretical Biology* **10**, 442–459 (1966).
- ⁵T. Stewart, L. Strijbosch, H. Moors, and P. van Batenburg, “A Simple Approximation to the Convolution of Gamma Distributions”, *SSRN Electronic Journal*, **10.2139/ssrn.900109** (2007).
- ⁶H. Paul, “Photon antibunching”, *Reviews of Modern Physics* **54**, 1061–1102 (1982).

- ⁷X. T. Zou and L. Mandel, “Photon-antibunching and sub-Poissonian photon statistics”, *Physical Review A* **41**, 475–476 (1990).
- ⁸D. L. Jones, R. C. Brewster, and R. Phillips, “Promoter architecture dictates cell-to-cell variability in gene expression”, *Science* **346**, 1533–1536 (2014).
- ⁹V. N. Kampen, “Stochastic processes in physics and chemistry”, in *Stoch. Process. Phys. Chem.* 3rd (2007).
- ¹⁰V. Shahrezaei and P. S. Swain, “Analytical distributions for stochastic gene expression”, *Proceedings of the National Academy of Sciences* **105**, 17256–17261 (2008).
- ¹¹M. Abramowitz and I. A. Stegun, eds., *Handbook of mathematical functions: with formulas, graphs, and mathematical tables* (1964).
- ¹²J. W. Pearson, S. Olver, and M. A. Porter, “Numerical methods for the computation of the confluent and Gauss hypergeometric functions”, *Numerical Algorithms* **74**, 821–866 (2017).
- ¹³A. Gil, J. Segura, and N. M. Temme, “Numerically satisfactory solutions of hypergeometric recursions”, *Mathematics of Computation* **76**, 1449–1469 (2007).
- ¹⁴J. Gunawardena, “Models in biology: ‘accurate descriptions of our pathetic thinking’”, *BMC biology* **12**, 29 (2014).
- ¹⁵A. Gelman, D. Simpson, and M. Betancourt, “The Prior Can Often Only Be Understood in the Context of the Likelihood”, *Entropy* **19**, 555 (2017).
- ¹⁶A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis*, Third edition, Chapman & Hall/CRC Texts in Statistical Science (CRC Press, Boca Raton, 2013).
- ¹⁷M. Betancourt, “A Conceptual Introduction to Hamiltonian Monte Carlo”, arXiv:1701.02434 [stat] (2018).
- ¹⁸B. Carpenter et al., “*STAN* : A Probabilistic Programming Language”, *Journal of Statistical Software* **76**, 10.18637/jss.v076.i01 (2017).
- ¹⁹S. Mahajan, *Street-fighting mathematics: the art of educated guessing and opportunistic problem solving* (MIT Press, Cambridge, Mass, 2010).
- ²⁰R. C. Brewster, D. L. Jones, and R. Phillips, “Tuning Promoter Strength through RNA Polymerase Binding Site Design in *Escherichia coli*”, *PLoS Computational Biology* **8**, edited by E. van Nimwegen, e1002811 (2012).
- ²¹H. G. Garcia and R. Phillips, “Quantitative dissection of the simple repression input-output function”, *Proceedings of the National Academy of Sciences* **108**, 12173–12178 (2011).

Appendix B

SUPPORTING INFORMATION FOR CHAPTER III, THEORETICAL INVESTIGATION OF A GENETIC SWITCH FOR METABOLIC ADAPTATION

B.1 Model choices, simplifications and assumptions

This part of supporting information text offers a more detailed discussion of the assumptions that were made in the model and explains the choices that were made.

Simplistic modeling of transcription. The precise mechanism at the promoter level is not the key feature of this system, and we expect the same qualitative results for slightly different transcription models. This work aims to gain an intuition from the results, and the more unknown parameters there are, the harder this becomes. Furthermore, the mechanistic details of transcription are not entirely known and would make the model too complex: for example, transcription probably is not a Poisson process, but comes in bursts. Additionally, very little is known about the function of the *xapAB* promoter. For these reasons, we do not model every detail, but focus on the most relevant features of the system, while still trying to be physically correct.

Firstly, the equilibrium assumption is not completely accurate, but has proven successful in previous work. An example for this is the very detailed analysis of a gene with one repressor in [1]. Secondly, different binding energies of XapR to its two binding sites would change the equations as follows: the factor $\left(\frac{1}{K_{\text{XapR}}}\right)^2$ would become $\frac{1}{K_{\text{XapR},1}K_{\text{XapR},2}}$ and the factor $\frac{2}{K_{\text{XapR}}}$ would be replaced by $\frac{K_{\text{XapR},1}+K_{\text{XapR},2}}{K_{\text{XapR},1}K_{\text{XapR},2}}$. For most choices of K_{XapR} , this does not give the exact same results quantitatively, but we do observe the same qualitative outcome (see also next paragraph). Thirdly, we consider all but one state as inactive, which is not strictly correct. Still, as mentioned in the main text, experiments suggest that this is reasonable. The fourth main simplification, which is the independence of polymerase and XapR binding, will be discussed in a separate subsection below.

In Fig B.1, we have plotted phase portraits where some of the above assumptions have not been made. For that, we have chosen some extreme values for the new parameters to demonstrate that these simplifications do not have a significant influence. The




Figure B.1: **Results for different models of transcription.** The parameters are the same as the ones that were used for all other phase portraits in this report and that are listed in a table in the main part. In (A), the phase portrait after addition of non-zero transcription rates to all other polymerase-bound states is shown (third assumption lifted). The new parameter values are $r_I = 10^{-2}r_p$ for the transcription rate when only one XapR is bound and $r_0 = 10^{-3}r_p$ for the basal rate. In (B), the phase portrait after introducing different XapR dissociation constants can be seen (second assumption lifted). The new parameter values are $[XapR]_{R,1} = 10$ and $[XapR]_{R,2} = 0.1$.

two additional transcription rates are large compared to experimental observations. Moreover, the chosen values of $[XapR]_{R,i}$ correspond to a difference of two orders of magnitude between $K_{XapR,1}$ and $K_{XapR,2}$. The difference between Oid and O3 in *lac* are two to three orders of magnitude [2]. The values were chosen such that $K_{XapR,1}K_{XapR,2} = K_{XapR}^2$. Unsurprisingly, this is necessary to not observe large deviations from the original equations. The results lead to the conclusion that, for the purposes of this paper, the aforementioned simplifications are appropriate.

Activator and polymerase interaction. A mechanistically more precise model of activation would include an interaction energy between polymerase and XapR. In that case, p_{active} cannot be factored into separate terms for the polymerase and the XapR binding probability. However, the precise activation mechanism is unknown.

Moreover, introducing an interaction energy between polymerase and XapR would make the equations overly complex. This would not only make factorization impossible, it would also be just as consequential to have different dissociation constants

Table B.1: **States and weights for a more precise model of transcription.**

XapR	Pol.	Transcription	Weight
none	no	0	1
first	no	0	w_{X1}
second	no	0	w_{X2}
both	no	0	$w_{X1}w_{X2}e^{\Delta\epsilon_{coop}}$
none	yes	r_0	w_p
first	yes	r_{X1}	$w_pw_{X1}e^{\Delta\epsilon_{X1}}$
second	yes	r_{X2}	$w_pw_{X2}e^{\Delta\epsilon_{X2}}$
both	yes	r_2	$w_pw_{X1}w_{X2}e^{\Delta\epsilon_{X1}+\Delta\epsilon_{X2}+\Delta\epsilon_{coop}}$

The first column shows which XapR sites are occupied in this state, the second one whether polymerase is bound to the promoter, and the third the transcription rate. In the weights, the notation is as follows: w_{X1} is the weight for XapR binding to the first binding site on the promoter, w_{X2} is that to the second site, w_p is the weight of polymerase binding to the promoter, $e^{\Delta\epsilon_{X1}}$ is the interaction energy between one XapR on the first binding site and the polymerase that is bound, $e^{\Delta\epsilon_{X2}}$ is the same for the second binding site, and finally, $e^{\Delta\epsilon_{coop}}$ is the interaction energy of the two XapR at both binding sites as in the main part.

and additional rates as in the previous paragraph. This leads to a huge number of terms and parameters because of different interaction energies for every state, which can be foreseen from the states and weights shown in the main text. The precise transcription mechanism is not one of the central parts of the circuit we are investigating (see also discussion in the previous paragraph). Thus, making it the by far most complex and precise part of the equations would be unreasonable and hinder the understanding of the more relevant components of the system.

Furthermore, we do not expect our approximation to have a large influence: the interaction energy will make the doubly bound state dominate, and the polymerase-XapR interaction should be weak in the singly bound state, because very weak activation is observed in the experiment when one of the XapR binding sites is removed (see section B.5 in this supporting information text). If $e^{\Delta\epsilon_{X1}}$ and $e^{\Delta\epsilon_{X2}}$ are small, we have $w_pw_{X1}e^{\Delta\epsilon_{X1}} \approx w_pw_{X2}e^{\Delta\epsilon_{X2}} \approx w_pw_X$. If, in addition, the transcription rates for the singly bound states are much smaller than those for the doubly bound state, we can neglect the singly bound states in the nominator of p_{active} (which, as argued in the previous paragraph “Simplistic modeling of transcription”, is a good assumption). Thereby, we can obtain the same form of p_{active} that we are working with in the main part by including the interaction between polymerase and XapR into ϵ_{coop} .

Working with only one protein species. The proteins are transcribed together, and thus, their rate of transcription and mRNA decay should be the same. The protein decay is assumed to be dominated by cell division, making the decay rates roughly equal as well. What could differ is the translation rate. Changes in this rate rescale the protein synthesis rate by the same factor and thus, in steady-state, change the total number of proteins. The only other place where the proteins appear in the ODE's is in their respective Michaelis-Menten term, and the latter is proportional to $k_{\text{cat}}[E]_0$. Hence, different numbers of XapA and XapB can be accounted for by changing their respective values of k_{cat} .

Exponential protein decay. The protein decay should be dominated by dilution through cell division, which is a discrete process. However, the cell grows before division, resulting in a continuous dilution in concentration. Furthermore, replication leads to a varying number of gene copies and thereby also a varying transcription rate. Trying to capture all these mechanistic details would lead to an overcomplicated model. Treating the protein decay as a Poisson process has produced correct results in past research, which is why we expect this assumption to be reasonable.

Kinetics of XapB, NupC, and NupG. All statements in this paragraph are based on data from [3]. Xanthosine entering the cell consists of two steps. In a first transport step, xanthosine passes the outer membrane through porins like Tsx, OmpF and OmpC. This does not seem to be a rate-limiting step. Then, XapB actively transports it across the inner membrane, powered by the proton motive force. There are two other active nucleoside transporters that seem to be transporting xanthosine with a very low affinity: NupC and NupG. It was found that $\Delta nupC \Delta nupG$ strains cannot grow on xanthosine. Hence, these seem to be necessary to “start” the system by importing small amounts of xanthosine that then activate *xapB*, and there appears to be no other significant way in which xanthosine can enter the cell.

As mentioned in the main text, we assume the kinetic scheme of all three transporters to be similar to that of *lac* permease (as it is described in [4]), which is discussed in more detail in the next subsection. One transporter of this type can do two things to change the intracellular xanthosine concentration $[x]$: it can actively transport one xanthosine into the cell, powered by the proton gradient, or it can transport one out of the cell, against the proton gradient. For net transport, a proton and a substrate need to bind on one side of the membrane and detach from the transporter on the other side of the membrane before the transporter goes back to the first side again (see next

subsection). Because of the proton gradient, influx is overall thermodynamically more likely than export at low intracellular xanthosine concentrations, which leads to a net influx. For much higher intra- than extracellular xanthosine concentrations, the difference in the chemical potential of xanthosine across the membrane can dominate that of the protons and there is net efflux.

We model influx and efflux separately and use Michaelis-Menten kinetics for XapB (this is discussed in more detail in the next subsection). For the turnover rate and Michaelis constant for influx we write $k_{b,i}$ and $K_{b,i}$, respectively, and $k_{b,e}$, $K_{b,e}$ for efflux. These parameters implicitly include the proton gradient.

For NupC and NupG, we approximate the transport reactions as linearly dependent on the substrate. This should be appropriate, since their affinity for xanthosine is so low that it cannot even be measured properly [3], which makes K_M much larger than c or $[x]$ will ever be. Thus, the general $k_{cat}[E]_0 \frac{[S]}{K_M + [S]}$ can be approximated as $\frac{k_{cat}}{K_M}[E]_0[S]$. In our equations, $\frac{k_{cat}}{K_M}[E]_0$ is denoted by k_{nup} for influx, and the efflux rate is written as ξk_{nup} . For the XapB kinetics, no such approximation can be made, because the xanthosine concentration in the dynamic system can range from far below the respective K_M value to far above.

The kinetics of *lac* permease. All transporters in our system are driven by the proton gradient [3], and we assume their kinetic scheme to be like that of *lac* permease. This scheme is shown schematically in Fig B.2. Even though it includes several steps, the essential part is substrate binding, conformational change of the transporter, and substrate release on the other side of the membrane. We assume the proton gradient across the membrane to be constant and proton binding and unbinding to the transporter to be fast. Hence, the latter does not need to be included explicitly in the kinetic model (it is implicitly part of the turnover rate).

Substrate release should be much faster than conformational change, which is analogous to ignoring the “EP” state in Michaelis-Menten (enzyme with product bound, here the conformationally changed transporter with substrate still bound). Also, it should be reasonable to neglect the immediate backwards reaction (before proton release), and thus, Michaelis-Menten kinetics is obtained. Data from [5] and [6] actually show that, at a physiologically reasonable proton gradient, both, influx and efflux, can be described by the Michaelis-Menten equation very well.

What is neglected in this modeling is the fact that the transporter needs to release the proton and undergo a conformational change again before being able to conduct the




Figure B.2: **Kinetic scheme for transport by proton-driven pumps.** The mechanism of lac permease as for example described in [4] is used as a model for xanthosine transport. The enzyme binds a proton and a substrate (steps 1 and 2), undergoes conformational change (steps 3 and 4), and then releases substrate and proton on the other side of the membrane (steps 5 and 6). When the enzyme is empty, it can change its conformation again to repeat the transport (steps 7 and 8). Instead, it could also bind a new proton and substrate and transport them in the opposite direction (steps 6, 5, 4, 3, 2, 1). Alternatively, it could also keep the proton and just cycle back and forth, exchanging substrates from outside and inside the cell (steps 5, 4, 3, 2).

same reaction again. However, by using experimentally determined apparent values to estimate the kinetic constants, this is already included in these parameter values. Another simplification we make is to treat influx and efflux as two independent reactions instead of thinking of it as one reversible kinetic process. This seems reasonable, though, because the transporter needs to lose its proton between reactions in order to obtain a net transport. If the transporter has for example done steps 1-6 in Fig B.2, it can from there on really enter two entirely different reactions: it can do its conformational change in its empty state (steps 7 and 8) to then transport a substrate again (steps 1-6) or it can bind a new proton and a new substrate directly and transport them back (steps 6-1). The former would give influx and the latter efflux, so the two of them can be modeled as separate reactions.

Passive transport across the membrane. Most small molecules diffuse passively through the membrane, but at physiological pH, xanthosine is mostly negatively charged and charged substances do not significantly diffuse across the membrane on their own. Nevertheless, from the pK_a of xanthosine we find that at physiological pH, about 1.5% of it is uncharged and could pass the membrane. This could be modeled as follows:

$$\left(\frac{d[x]}{dt} \right)_{\text{diff}} = JA = p(c - [x])A = \underbrace{pA}_{\text{rate}} (c - [x]) \quad (\text{B.1})$$

Here, J is the flux through the membrane, which follows directly from Fick's first law (if steady-state diffusion is assumed). The other new parameters are the membrane area A and the membrane permeability coefficient p of the diffusing substrate.

The diffusion rate $k_d = pA$ can be estimated with the help of Overton's rule. In our model, we find that this passive diffusion alone (i.e. without NupC and NupG) is enough to obtain bistability in the deterministic analysis as well as bimodality in the stochastic simulations. However, experiments show that $\Delta nupC \Delta nupG$ strains cannot grow on xanthosine [3]. This suggests that passive diffusion alone is not enough to make the cells switch. Experiments with cells in minimal media containing some glucose and a xanthosine concentration in the bimodal regime show that the switching time can be several generations. The same is observed in our simulations. The xanthosine concentration that was used for the double knockout experiments, where no glucose was present, is just above the switching threshold, which is when switching is slow. For this reason, we suspect that without NupC and NupG and with only xanthosine as an energy resource, the cells starve before they have switched. This probably is the reason why passive diffusion is, in reality, not enough to observe switching. NupC and NupG can probably keep the cells going at very low intracellular xanthosine concentrations. Because of the active transport, they can accumulate a higher intracellular than extracellular concentration of xanthosine.

Since NupC and NupG transport seems to be much more significant than diffusion, we do not include the latter explicitly in our model. However, because it takes the form given in Eq B.1, it can easily be accounted for implicitly by slightly changing the values of k_{nup} and ξ .

Other mechanisms for xanthosine degradation. There probably are some ways in which xanthosine is degraded, other than by XapA. Otherwise, initially uninduced

cells that are transferred to only xanthosine as an energy source might not survive until they have switched. As mentioned previously, it was found that they do [3], so they need to be able to metabolize xanthosine somehow. Nevertheless, these degradation pathways cannot have a big effect: as argued later on when estimating the Nup parameters, the intracellular xanthosine concentration in the uninduced state must be a bit larger, but not too much larger, than the extracellular one. Thus, degradation of xanthosine has to be small because import through Nup is weak. Due to this, it should be appropriate to linearize the equations describing these effects. Then, they can simply be accounted for in the term for Nup transport by making slight changes in the values of ξ and k_{nup} . It could, however, also be that the basal level of XapA is enough to obtain the small degradation effect without any other pathways.

B.2 Parameter estimation

In this section we present a detailed estimation of the values of all nondimensional parameters. Two dimensionful parameters appear in most nondimensional ones, namely γ_p and K_a . For this reason, we start by estimating these.

The protein decay rate γ_p : Protein decay is generally dominated by dilution through cell growth and division. The average doubling time of *E. coli* is 20-30 minutes. We take this as an estimate for the protein half life, which leads us to a decay rate of $\gamma_p = \frac{\ln 2}{t_{1/2}} \approx 5^{+2}_{-3} \cdot 10^{-4} \text{ s}^{-1}$. Note that these upper and lower bounds correspond to half-lives between 15 and 60 minutes, which should be reasonable. For the following nondimensionalization, we work with a fixed value of $\gamma_p = 5 \cdot 10^{-4} \text{ s}^{-1}$.

The Michaelis-Menten parameters for XapA (K_a , k_a): The Michaelis-Menten parameters of XapA kinetics have been measured in two independent publications [7, 8]. Both have found similar values and we conclude $K_a \approx 6 \pm 3 \cdot 10^{-5} \text{ M}$ and $k_a \approx 10^{-1 \pm 0.8} \text{ s}^{-1}$. This gives $k_a \approx 10^{2 \pm 0.8}$. We choose to work with $K_a = 5 \cdot 10^{-5} \text{ M}$ in the following.

The extracellular xanthosine concentration $[c]_a$: In the experiment, switching occurred if cells were grown in a solution with a concentration of xanthosine of roughly a few mM [3]. According to these experimental observations, the interesting range of xanthosine should be $[c]_a \in [0, 10^3]$ (nondimensionalized), which

corresponds to $c \in [0, 5 \cdot 10^{-2}]$ M. However, we are careful because there are issues with the solubility of xanthosine.

The MWC parameters $K_{\chi A}$, K_{IA} , and $\Delta\epsilon_x$: The three MWC parameters are the energy barrier between the active and inactive state of XapR and the equilibrium constants of xanthosine binding to XapR in the two states. In Fig B.3, the induction curve as a function of the nondimensional parameters is shown. We assume that evolution has tuned the induction curve such that less than 10% of XapR is active in the absence of xanthosine, and more than 90% is active for $[x]$ approaching infinity. This is a relatively weak assumption, otherwise XapR would not really be fulfilling its job. It helps us by putting a constraint on K_{IA} and $\Delta\epsilon_x$:

$$p_{\text{active}}([x]_a = 0) = \frac{1}{1 + e^{\Delta\epsilon_x}} < 0.1$$

$$1 - p_{\text{active}}([x]_a \rightarrow \infty) = \frac{e^{\Delta\epsilon_x}}{K_{IA}^2 + e^{\Delta\epsilon_x}} < 0.1$$

Thus,

$$\Delta\epsilon_x > 2.2 \quad \text{and} \quad K_{IA} > 3e^{\Delta\epsilon_x/2}.$$

A further constraint is found from $K_{IA} = \frac{K_{xI}}{K_{xA}} = e^{\beta\Delta E}$, where ΔE is the energy difference between xanthosine binding to the two states. It is unlikely that this energy difference is more than $7 k_B T$, which corresponds to roughly three strong hydrogen bonds. This gives us an upper bound $K_{IA} < 10^3$ and thereby $\Delta\epsilon_x < 10^1$. For $K_{IA} = 10^2$, the upper bound is $\Delta\epsilon_x < 7$. This gives us a range of $K_{IA} \approx 10^1 - 10^3$ and $\Delta\epsilon_x \approx 2 - 10$ or, more precisely, $\Delta\epsilon_x \approx 2$ to $\ln\left(\frac{1}{9} (K_{IA})^2\right) \approx 2(\ln(K_{IA}) - 1) < 12$. Compared with other systems, these values do seem reasonable.

To find $K_{\chi A}$ we furthermore assume that the steep part of the induction curve lies in the biologically relevant xanthosine regime, which estimated experimentally, see the paragraph about $[c]_a$ above. Again, this is a rather weak assumption. As can be seen from Fig B.3, this assumption sets the value range of $K_{\chi A}$. The line for $\Delta\epsilon_x = 5$ in Fig B.3 is at the experimentally expected position in the xanthosine regime. It becomes clear from the figure that the position of the steep range is strongly affected by $\Delta\epsilon_x$, and a change of 1 in the latter leads to roughly a change of 10^1 in the former. Thus, we estimate $K_{\chi A} \approx 10^{2 \pm 1} \cdot 10^{\Delta\epsilon_x - 5}$ and thereby finish our estimation of the MWC parameters.




Figure B.3: **Induction curves given by MWC model.** The main plot shows the probability of XapR being in the active state as a function of the logarithm of the xanthosine concentration for $\Delta\epsilon_x$ as shown, $K_{\chi A} = 10^2$, and $K_{IA} = 10^2$. The shaded areas indicate where the curve lies for changes in K_{IA} of one order of magnitude. Changing $K_{\chi A}$ by one order of magnitude, on the other hand, shifts all the curves by one order of magnitude in the corresponding direction to the left or right. The inset shows how the allowed regimes of $\Delta\epsilon_x$ and K_{IA} depend on each other. This explains why for $\Delta\epsilon_x = 7$ and $K_{IA} = 10$ (lower green line in the main plot), p_{active} at infinite xanthosine concentrations becomes far too low.

The XapR concentration $[\text{XapR}]_R$: It is experimentally known that the XapR copy number is very low, on the order of 10^{-8} M (10^1 molecules per cell) [9]. We estimate the XapR – DNA dissociation constant from the values for *lac* that were found in [2] and obtain $10^{-8\pm 2}$ M, where we have added one additional order of magnitude as insecurity to the upper and lower bound. This gives $[\text{XapR}]_R = \frac{[\text{XapR}]_{\text{tot}}}{K_{\text{XapR}}} \approx 10^{0\pm 2}$.

The XapR cooperative energy $\Delta\epsilon_{\text{coop}}$: Any cooperative energy should come mainly from an interaction between the two XapR molecules. It would be surprising if it was larger than $10 k_B T$, which corresponds to roughly 4 strong hydrogen bonds. Hence, we estimate $\Delta\epsilon_{\text{coop}} \approx 0 - 10$.

The transcription rate parameter ρ_m : The transcription rate $r_m \frac{[P]}{[P] + K_P}$ in the presence of two activator molecules should be on the order of $10^{-2\pm 1}$ nM s $^{-1}$, thus $\rho_m := \frac{r_m}{\gamma_p K_a} \frac{[P]}{K_P + [P]} \approx 10^{-3\pm 2}$. Note that the units are nM s $^{-1}$ and not just s $^{-1}$ because

the corresponding rate equation is of the form $\frac{d[m]}{dt} = r \cdot p$ where the probability p is dimensionless and the left-hand side has dimensions of mRNA concentration per time.

This estimate comes from the assumption that this rate should be roughly $10^1 - 10^2$ times larger than a normal basal rate (of other genes), for which common values are around $10^{-3 \pm 1} \text{ nM s}^{-1}$ [10]. Because experiments show that compared to the activated promoter, the basal transcription rate is almost zero for *xapAB* (probably due to the weak polymerase binding site), this actually means a 10^2 - to 10^3 -fold increase in the transcription rate due to activation by XapR.

Keep in mind that the rate includes not only the effect of XapR on the pure speed of transcription, but also any effect of XapR on the polymerase dissociation constant. Still, we do not expect higher values such as 10^0 nM s^{-1} since this is roughly the rate for transcription of rRNA and the latter should be significantly faster. To further validate our estimate, we can compare it to measured values for *lac*, which should be roughly the same because the gene function is similar. The transcription rate for *lac* in the absence of any repression that was found in [11] corresponds well with our estimate.

The mRNA decay rate parameter γ_{mp} . The mRNA life time is usually on the order of 1 to 8 minutes which gives a decay rate of $\gamma_m \approx 10^{-3} - 10^{-2} \text{ s}^{-1}$. Thus, we have $\gamma_{mp} := \frac{\gamma_m}{\gamma_p} \approx 10^{1 \pm 0.5}$.

The translation rate parameter ρ_p . The translation rate should be on the order of $r_p \approx 10^{-2} - 10^{-1} \text{ s}^{-1}$ [10]. This gives $\rho_p := \frac{r_p}{\gamma_p} \approx 10^{2 \pm 0.5}$.

The Michaelis-Menten parameters for XapB ($k_{\beta,i}, K_{\beta,i}, k_{\beta,e}, K_{\beta,e}$): To estimate the parameters for both influx and efflux, we need data from experiments where these two processes have been measured independently. Lactose permease, which we expect to function in a similar way as XapB, is comparatively well studied such that these rates have been measured in multiple ways and for different conditions. Some data from several publications is collected in [6] and we will base our estimates on this and the corresponding references.

Comparing the effective rates and Michaelis constants of lactose permease and the nucleoside transporters NupC and NupG shows that their orders of magnitude are similar: from [12, 13] we find $k_{cat} \approx 10^1 \text{ s}^{-1}$, $K_M \approx 10^1 - 10^2 \text{ } \mu\text{M}$ for LacY and

from [3, 14, 15] $k_{\text{cat}} \approx 10^2 \text{ s}^{-1}$, $K_{\text{M}} \approx 10^0 - 10^1 \mu\text{M}$ for pyrimidines through NupC and NupG. Note that especially for Nup, these are only rough measurements, and it is known that NupC behaves different with purines (thus also with xanthosine). Given these values, we expect lactose permease to be similar enough to nucleoside transporters so that we can use it as a starting point for our estimation of the XapB kinetic values.

We use the measured influx values that are summarized in [6] for a membrane potential of 100 mV and obtain $k_{\text{b,i}} \approx 10^{1\pm1} \text{ s}^{-1}$ and $K_{\text{b,i}} \approx 10^{-4\pm2} \text{ M}$. Here, the uncertainty is estimated from the distribution of values of many kinds of enzymes that is presented in [10]. The bounds also reflect the difference between LacY and Nup that was found above. For the dimensionless quantities, this gives $k_{\beta,i} \approx 10^{4\pm1}$ and $K_{\beta,i} \approx 10^{1\pm2}$.

For the efflux parameters, we proceed similarly and obtain $k_{\text{b,e}} \approx 10^{0\pm2} \text{ s}^{-1}$ and $K_{\text{b,e}} \approx 10^{-3\pm3} \text{ M}$, therefore $k_{\beta,e} \approx 10^{3\pm2}$ and $K_{\beta,e} \approx 10^{2\pm3}$. Because of disagreement in literature, difficulty in measuring these parameters, and a lack of knowledge about what value range should be expected, we put larger error bounds than for the influx.

The non-specific transport parameters k_{η} , ξ : Estimating the rates of the Nup transporters for xanthosine is difficult since there are no experimental values that could be used as a reference. Because the rate cannot even be measured [3], it must be much lower than that of XapB. The kinetic parameters should, however, be such that the steady-state intracellular xanthosine concentration is significantly larger than the extracellular one. Otherwise, there would be no relevant difference to passive diffusion, which was shown to be insufficient for switching [3]. This essentially means $\xi < 1$, which on the other hand implies that the Nup rate is larger than or at least equal to the diffusion rate: If Nup transport were slower than diffusion, the latter would rule and ξ would be close to 1 (because this is the ξ value of diffusion).

In addition, we know from *lac* permease that influx is not much faster than efflux, which implies that ξ also cannot be too much smaller than 1. We guess $\xi \approx 0.8 \pm 0.1$. The number of Nup transporters per cell is roughly $10^2 - 10^3$ [9]. We expect $\frac{k_{\text{cat}}}{K_{\text{M}}}$ to be at least two or three orders of magnitude lower for Nup than for saturated XapB, where we have $\frac{k_{\text{b}}}{K_{\text{b}}} \approx 10^{5\pm3} \text{ M}^{-1} \text{ s}^{-1}$. These arguments set an estimated upper bound on the non-specific transport rate: $k_{\text{nup}} = \frac{k_{\text{cat}}}{K_{\text{M}}} [\text{E}]_0 \lesssim 10^{-3\pm3} \text{ s}^{-1}$.

To find a lower bound, we estimate the diffusion rate. As described in Eq B.1 in the previous section, the diffusion rate can be written as $10^{-2} \cdot p \cdot A$, where p

is the permeability coefficient of the membrane for xanthosine, A is the total cell surface area, and 10^{-2} is the fraction of uncharged xanthosine at physiological pH. From the permeabilities of various substances that are collected in [10] and from the measured permeability values for purine nucleosides in [16], we estimate a permeability coefficient of $\tilde{p} = 10^{-8 \pm 3} \text{ m s}^{-1}$. We write the tilde because this value is given in a convention where $\tilde{p} \cdot A = \frac{dx}{dt}$ with x being the number of xanthosine molecules and not the concentration. Since we work in concentrations, we divide by an estimated cell volume of 10^{-18} m^3 and, with a cell surface area of 10^{-11} m^2 , find a diffusion rate of roughly $10^{-3 \pm 3} \text{ s}^{-1}$. This sets the lower bound of k_{nup} .

Interestingly, the upper and lower bound turn out to be the same. Thus, we obtain the range $k_{\text{nup}} \approx 10^{-3 \pm 3} \text{ s}^{-1}$, meaning $k_{\eta} \approx 10^{-1 \pm 3}$.

B.3 Additional plots and explanations of the results

The 2D set of equations. If the mRNA concentration is set to steady-state as discussed in the main text, the dynamical system is reduced to the two equations shown below:

$$\frac{d[p]_a}{d\tau} = \rho \frac{[\text{XapR}]_{R,A}^2 e^{-\Delta\epsilon_{\text{coop}}}}{1 + 2[\text{XapR}]_{R,A} + [\text{XapR}]_{R,A}^2 e^{-\Delta\epsilon_{\text{coop}}}} - [p]_a \quad (\text{B.2})$$

$$\frac{d[x]_a}{d\tau} = \left(k_{\beta,i} \frac{[c]_a}{K_{\beta,i} + [c]_a} - k_{\beta,e} \frac{[x]_a}{K_{\beta,e} + [x]_a} - k_{\alpha} \frac{[x]_a}{1 + [x]_a} \right) [p]_a + k_{\eta} ([c]_a - \xi[x]_a) \quad (\text{B.3})$$

$$\text{with } [\text{XapR}]_{R,A} = [\text{XapR}]_R \frac{\left(1 + \frac{[x]_a}{K_{\chi A}}\right)^2}{\left(1 + \frac{[x]_a}{K_{\chi A}}\right)^2 + e^{\Delta\epsilon_{\chi}} \left(1 + \frac{[x]_a}{K_{\chi A}} \frac{1}{K_{IA}}\right)^2}.$$

Here, several of the parameters from the 3D system could now be combined into one by defining $\rho := \frac{\rho_m \rho_p}{\gamma_{\text{mp}}}$.

Magnitude of flow in the phase portraits. The plots in the main part do not show the magnitude of the vector fields. For additional information, we show a 2D vector plot with the magnitudes in Fig B.4.

Increasing $[c]_a$. Fig B.5 shows a family of nullclines for increasing extracellular xanthosine concentration. It shows the transition from monostability, to bistability, and back to monostability again.




Figure B.4: **Vector plot of a standard case of bistability.** The 2D phase portraits as a vector plot for the standard set of parameters that was used in the main text.




Figure B.5: **A family of curves for increasing $[c]_a$.** The blue, dashed curve is the protein nullcline, which does not change with $[c]_a$. The other curves are the xanthosine nullclines for different values of the extracellular xanthosine concentration.

Including transcriptional bursting. To test the robustness of our results to larger stochasticity (which is expected in the real biological system), we explored how the addition of burstiness in transcription altered our stochastic simulations. We assumed a geometric distribution for the burst size s , i.e. $p(s|\theta) = \theta \cdot (1 - \theta)^{s-1}$ with mean $\theta^{-1} = 4$, which is comparable to or slightly larger than reported values [17]. We then truncate this distribution at a burst size of 10, meaning we did not consider larger bursts, leading to an actual mean of around 3.2. The transcription rate was rescaled such that the mean number of mRNA's produced remained the same.

We observe the same qualitative behavior but with stronger fluctuations around the mean and larger variation in the time to reach the steady-state mean. As an example, Fig B.6A shows the time evolution of one run of the simulation with bursting for the same conditions (i.e. parameters and starting values) as the time evolution that was shown in the main text. In addition, the distribution of adaptation times can be seen in Fig B.6B. Fig B.7 shows the distribution obtained for the same and an additional condition (i.e. parameters and starting values) as for the distributions shown in the main text. The variance in the timescales and distributions is larger than without bursting. Furthermore, switching already happens for lower $[c]_a$: in B.7C, most cells are in the high expression state, whereas in the corresponding distribution in the main text, there were many uninduced cells, i.e. the first peak was higher. To obtain a bimodal distribution like in the main text, $[c]_a$ now needs to be set to 17.5 instead of 18.5 as shown in B.7B. This was to be expected, since the larger fluctuations make it easier for the system to jump across the middle fixed point.

Still, there are no big qualitative differences to the results without bursting. Thus, the assumptions we made that lead to less stochasticity than in the real biological system are not expected to have a very significant influence on our results, and the switch-like behavior seems to be stable to stochastic perturbations.

Hysteresis effects. In Fig B.8, the results mentioned in the hysteresis part of the main text are shown. As mentioned there, the simulation was started with initially fully induced cells. At concentrations where no cells switched from the low to the high expression state before, cells here do not switch back to low expression. The extracellular xanthosine concentration (for the other parameters being as in the Table in the main text) needed to be as low as $[c]_a = 10$ to observe bimodality and $[c]_a = 9$ to have all cells in the low expression state (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M).

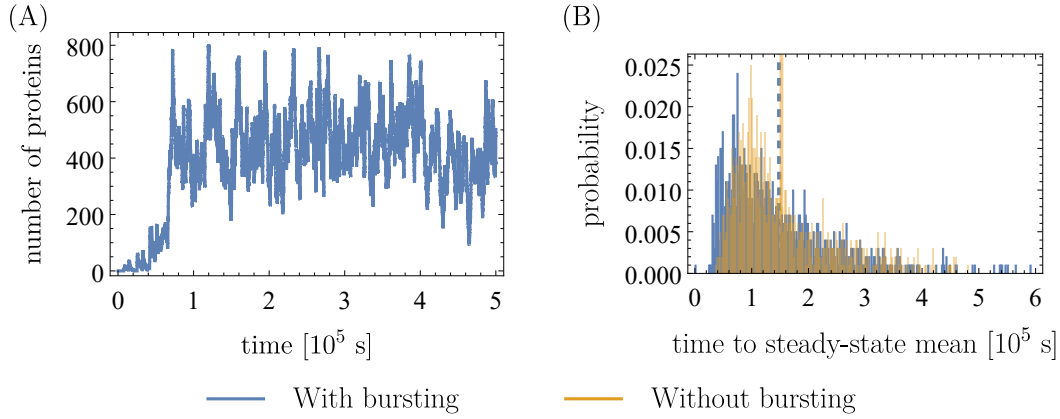


Figure B.6: Time evolution and adaptation times from the stochastic simulation when transcriptional bursting is included. (A) shows the time evolution from one run of the simulation. The simulation was run once for $5 \cdot 10^5$ s and started at an mRNA, protein and intracellular xanthosine count of 0. The parameters that were used are the same as in the table in the main text, the only exception being the extracellular xanthosine concentration, which was chosen to be $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M). (B) Blue shows the time until 90% of the protein concentration at the upper fixed point was reached in 1000 runs of the simulation (same conditions as (A)). The distribution from the main text (without bursting) is shown in orange for comparison. For both, the mean of the distribution is indicated by the line in the corresponding color.

B.4 Chemical master equation

Here we provide some further details on our stochastic simulations by writing the chemical master equation for our model. The state of our system is described by the number of mRNA m , protein p , and xanthosine molecules x , and we write $P(m, p, x)$ for the probability of that state. The chemical master equation for the

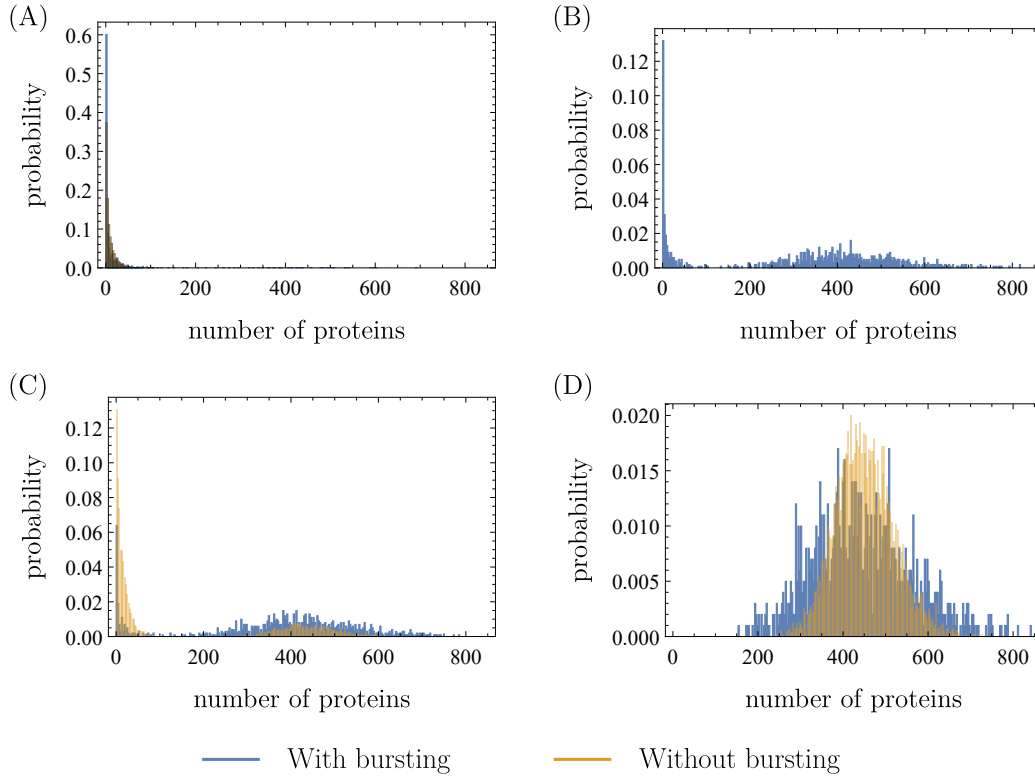


Figure B.7: Distributions from the stochastic simulation when transcriptional bursting is included. Blue shows the distributions from the simulation with bursts, orange shows those without bursts from the main text for comparison. The following values of $[c]_a$ were used: (A) $[c]_a = 12$, (B) $[c]_a = 17.5$ (not shown in main text), (C) $[c]_a = 18.5$, and (D) $[c]_a = 25$ (recall $[c]_a := \frac{c}{K_a}$, $K_a = 5 \cdot 10^{-5}$ M). All other parameters are as in the table in the main text, and the simulation was run 1000 times for 10^6 s each (simulated time) and started at a mRNA, protein and intracellular xanthosine count of 0.

system is:

$$\begin{aligned}
 \frac{dP(m, p, x)}{dt} = & -P(m, p, x)[r_m p_{\text{active}}(x) + \gamma_m m + r_p m + \gamma_p p \\
 & + k_{b,i} \frac{c}{K_{b,i} + c} p + k_{b,e} \frac{x}{K_{b,e} + x} p + k_a \frac{x}{K_a + x} p \\
 & + k_{\text{nup}} c + k_{\text{nup}} \xi x] \\
 & + P(m+1, p, x) \gamma_m (m+1) \\
 & + P(m-1, p, x) r_m p_{\text{active}}(x) \\
 & + P(m, p+1, x) \gamma_p (p+1) \\
 & + P(m, p-1, x) r_p m \\
 & + P(m, p, x+1) [k_{b,e} \frac{x+1}{K_{b,e} + x+1} p \\
 & + k_a \frac{x+1}{K_a + x+1} p + k_{\text{nup}} \xi (x+1)] \\
 & + P(m, p, x-1) [k_{b,i} \frac{c}{K_{b,i} + c} p + k_{\text{nup}} c]
 \end{aligned} \tag{B.4}$$

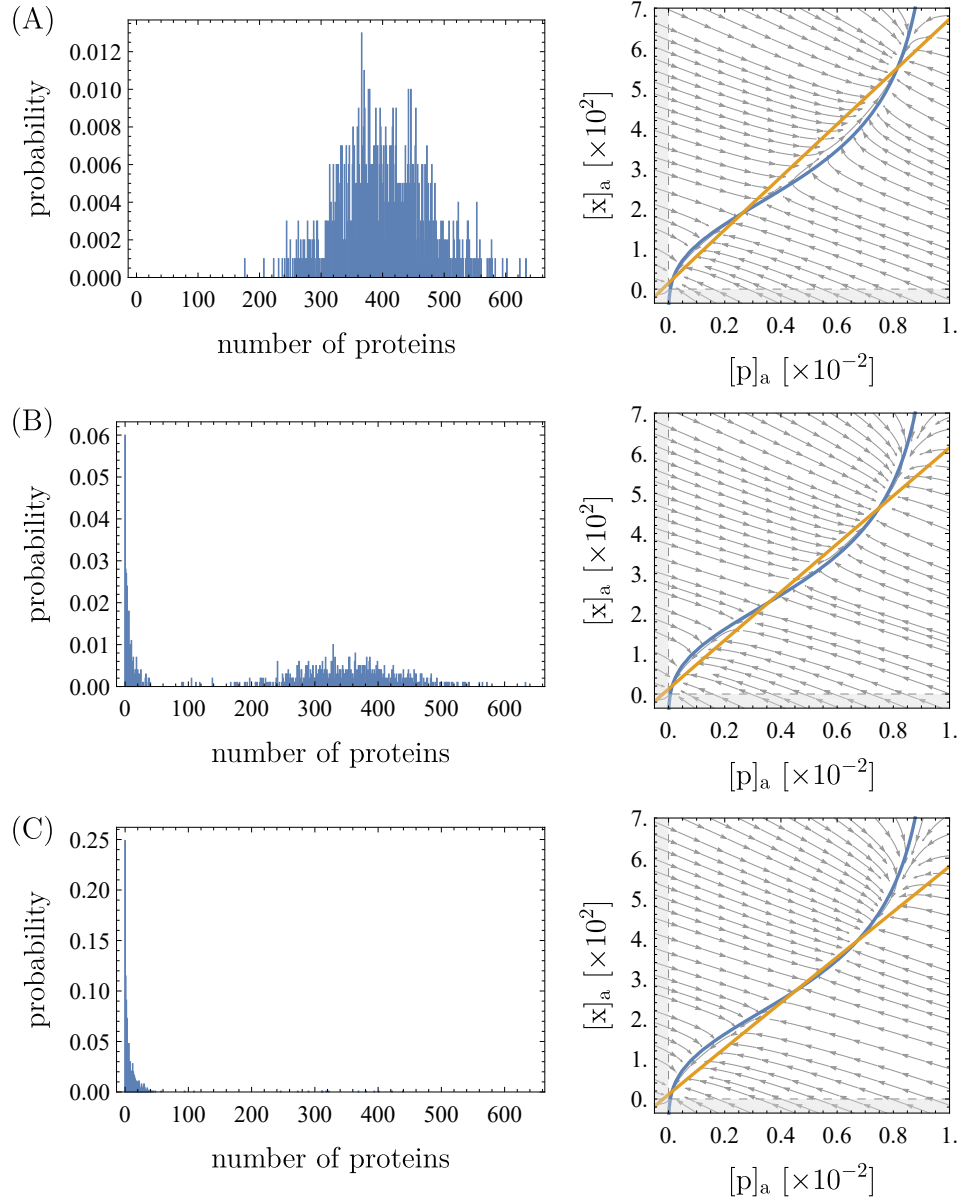


Figure B.8: **Decreasing $[c]_a$ in the fully induced system.** Apart from $[c]_a$, the parameters are the same as in the table in the main text. For the distributions, the simulations were run 1000 times for 10^6 s each and started at the mRNA, protein and intracellular xanthosine counts of the high fixed point in the corresponding phase portraits. The values of $[c]_a$ are 12 in (A), 10 in (B), and 9 in (C). When comparing to the distributions in the main text where the simulation had initial concentrations of 0, a clear hysteresis effect can be observed.

Here, all the parameters of course have to be measured in counts, not in concentrations. As can be seen from the equation, none of the propensities depends explicitly on time and hence, we can use the hybrid algorithm between classical Gillespie and τ -leaping that we relied on.

Note that in the case with bursts that is discussed above, the transcription term $r_{\text{m}}p_{\text{active}}(x)P(m-1, p, x)$ is changed to $r_{\text{m}}p_{\text{active}}(x) \sum_{k=1}^{10} P(m-k, p, x) \frac{\theta \cdot (1-\theta)^{k-1}}{Z}$ with $Z = \sum_{k=1}^{10} k \cdot \theta \cdot (1-\theta)^{k-1}$ and $\theta = 1/4$ in our example.

B.5 Experimental materials and methods

Construction of $\Delta xapABR$ deletion strain. Cells containing plasmid pSIM6 were diluted 1:100 in 50 mL LB media and grown to an OD600 of ≈ 0.4 at 30°C. The culture was immediately placed in a water bath shaker at 43 °C for 15 minutes and then cooled in an ice bath for 10 minutes. Cells were then spun down for 10 minutes (4,000 g, 4 °C) and resuspended on ice in 50 mL of chilled water. This was repeated three times before resuspending in 200 μ L of chilled water to generate competent cells. Homologous primer extension sequences were obtained from Baba *et al.* [18] and used to generate linear DNA containing a kanamycin resistance gene insert by PCR (amplified from plasmid pKD4), using primers that contained homology for the region on the chromosome to be deleted [19]. Electroporation was performed using 1 μ L purified PCR product (about 100 ng DNA), mixed with 50 μ L cells. Cells were resuspended in 750 μ L SOC media and outgrowth was performed on a shaker at 30 °C for 90-120 minutes. Cells were then selected on an LB-agar plate containing kanamycin (30 μ g/mL) and grown overnight at 30 °C. The deletions were confirmed by both colony PCR and sequencing. After confirmation, the deletion was transferred to a clean strain through P1 transduction and selection on kanamycin. Removal of kanamycin resistance was achieved by transforming cells with plasmid pCP20 (expressing FLP recombinase, carbenicillin resistance, on a temperature-sensitive replicon), with cells grown at 30 °C. Lastly, several colonies were then growth colony-purified once non-selectively at 43 °C and then tested for loss of antibiotic resistance for both kanamycin and carbenicillin.

Construction of $\Delta xapAB$ strain. Further experiments are underway in which precise control of the xapR copy number are needed and the preliminary data from these experiments are shown in the right panel of Fig. 2A of the main text. This strain was constructed by creating a XapR-mCherry fusion protein expressed from a tetR regulated promoter. This fusion construct was integrated into the ycbN locus

Primer	Sequence	Comments
<i>xapABR</i> <> P2	actatttcggcaattcctgtctctcac ctactgtgtcaatgcagccaa- CATAT GAATATCCTCCT- TAG	Forward primer for pKD4 with homology to <i>xapABR</i> locus
<i>xapABR</i> <> P1	agatacaacaggtatcgctaataatagagg aaagtgtcagagttaaagat- gtttt ttcacGTGTAGGCTG- GAGCTGCTTC	Reverse primer for pKD4 with homology to <i>xapABR</i> locus
<i>xapABR</i> -seq-fwd	cgggtcgtagctcagttggta	upstream of <i>xapR</i> , to check <i>xapABR</i> deletion
<i>xapABR</i> -seq-rev	ttccagtcgaagaacccgtac	downstream of <i>xapA</i> , to check <i>xapABR</i> deletion

Table B.2: **Primers for deletion strains.** Lower case sequences are homologous to a chromosomal locus. Upper case indicates sequences used for PCR amplification.

of the $\Delta xapABR$ strain making an effective $\Delta xapAB$ strain. A pZS3-PN25-TetR expression plasmid was also inserted into this genetic background which permits titratable control over the *xapR*-mCherry copy number. For the data presented in this work, expression was maximized using 10 ng/mL of the TetR inducer anhydro-tetracycline. Further information of a similar construct (LacI-mCherry) can be found in [20] and [21].

References

- ¹R. Phillips, N. M. Belliveau, G. Chure, H. G. Garcia, M. Razo-Mejia, and C. Scholes, “Figure 1 Theory Meets Figure 2 Experiments in the Study of Gene Expression”, *Annual Review of Biophysics* **48**, 121–163 (2019).
- ²M. Razo-Mejia, S. L. Barnes, N. M. Belliveau, G. Chure, T. Einav, M. Lewis, and R. Phillips, “Tuning Transcriptional Regulation through Signaling: A Predictive Theory of Allosteric Induction”, *Cell Systems* **6**, 456–469.e10 (2018).
- ³M. H. H. Nørholm and G. Dandanell, “Specificity and Topology of the Escherichia coli Xanthosine Permease, a Representative of the NHS Subfamily of the Major Facilitator Superfamily”, *Journal of Bacteriology* **183**, 4900–4904 (2001).
- ⁴H. R. Kaback, “A chemiosmotic mechanism of symport”, *Proceedings of the National Academy of Sciences* **112**, 1259–1264 (2015).
- ⁵G. J. Kaczorowski, D. E. Robertson, and H. R. Kaback, “Mechanism of lactose translocation in membrane vesicles from Escherichia coli. 2. Effect of imposed $\Delta \Psi$, ΔpH , and $\Delta \mu_{\text{H}^+}$ ”, *Biochemistry* **18**, 3697–3704 (1979).

- ⁶P. Viitanen, M. L. Garcia, and H. R. Kaback, “Purified reconstituted lac carrier protein from *Escherichia coli* is fully functional.”, *Proceedings of the National Academy of Sciences* **81**, 1629–1633 (1984).
- ⁷G. Dandanell, R. H. Szczepanowski, B. Kierdaszuk, D. Shugar, and M. Bochtler, “*Escherichia coli* Purine Nucleoside Phosphorylase II, the Product of the *xapA* Gene”, *Journal of Molecular Biology* **348**, 113–125 (2005).
- ⁸G. W. Koszalka, J. Vanhooke, S. A. Short, and W. W. Hall, “Purification and properties of inosine-guanosine phosphorylase from *Escherichia coli* K-12.”, *Journal of Bacteriology* **170**, 3493–3498 (1988).
- ⁹G.-W. Li, D. Burkhardt, C. Gross, and J. S. Weissman, “Quantifying Absolute Protein Synthesis Rates Reveals Principles Underlying Allocation of Cellular Resources”, *Cell* **157**, 624–635 (2014).
- ¹⁰R. Milo and R. Phillips, *Cell biology by the numbers* (Garland Science, Taylor & Francis Group, New York, NY, 2016), 356 pp.
- ¹¹M. Razo-Mejia, S. Marzen, G. Chure, R. Taubman, M. Morrison, and R. Phillips, “First-principles prediction of the information processing capacity of a simple genetic circuit”, arXiv:2005.03214 [q-bio] (2020).
- ¹²K. Dornmair, P. Overath, and F. Jähnig, “Fast measurement of galactoside transport by lactose permease”, *The Journal of Biological Chemistry* **264**, 342–346 (1989).
- ¹³J. K. Wright and R. Seckler, “The lactose/H⁺ carrier of *Escherichia coli*: lac YUN mutation decreases the rate of active transport and mimics an energy-uncoupled phenotype”, *Biochemical Journal* **227**, 287–297 (1985).
- ¹⁴Y. Komatsu and K. Tanaka, “Deoxycytidine uptake by isolated membran vesicles from *Escherichia coli* K 12”, *Biochimica et Biophysica Acta (BBA) - Biomembranes* **311**, 496–506 (1973).
- ¹⁵A. Munch-Petersen, B. Mygind, A. Nicolaisen, and N. J. Pihl, “Nucleoside transport in cells and membrane vesicles from *Escherichia coli* K12”, *The Journal of Biological Chemistry* **254**, 3730–3737 (1979).
- ¹⁶C. Wei and A. Pohorille, “Permeation of Nucleosides through Lipid Bilayers”, *The Journal of Physical Chemistry B* **115**, 3681–3688 (2011).
- ¹⁷I. Golding, J. Paulsson, S. M. Zawilski, and E. C. Cox, “Real-Time Kinetics of Gene Activity in Individual Bacteria”, *Cell* **123**, 1025–1036 (2005).
- ¹⁸T. Baba et al., “Construction of *Escherichia coli* K-12 in-frame, single-gene knock-out mutants: the Keio collection”, *Molecular Systems Biology* **2**, 10 . 1038 / msb4100050 (2006).
- ¹⁹K. A. Datsenko and B. L. Wanner, “One-step inactivation of chromosomal genes in *Escherichia coli* K-12 using PCR products”, *Proceedings of the National Academy of Sciences* **97**, 6640–6645 (2000).

- ²⁰R. C. Brewster, F. M. Weinert, H. G. Garcia, D. Song, M. Rydenfelt, and R. Phillips, “The Transcription Factor Titration Effect Dictates Level of Gene Expression”, *Cell* **156**, 1312–1323 (2014).
- ²¹G. Chure, Z. A. Kaczmarek, and R. Phillips, “Physiological Adaptability and Parametric Versatility in a Simple Genetic Circuit”, *bioRxiv*, 10.1101/2019.12.19.878462 (2019).

