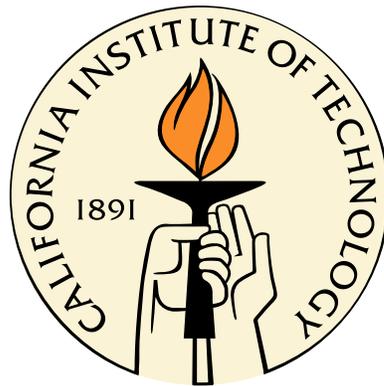


Reduced Order Models for Open Quantum Systems

Thesis by
Asa Sies Hopkins

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



California Institute of Technology
Pasadena, California

2009
(Defended November 17, 2008)

© 2009

Asa Sies Hopkins

Some Rights Reserved

Creative Commons Attribution 3.0 United States

Acknowledgments

Many thanks to my advisor, Hideo Mabuchi, for giving me the freedom and flexibility to pursue this project. I worked on several different projects during my graduate studies, and I appreciate Hideo for always standing behind my choices, and always being ready with alternatives, suggestions, and support whenever I asked.

The research presented in this thesis would not have been possible without the help of Ramon van Handel and Luc Bouten. Ramon set a terribly high example for the quality of work to be achieved, and yet was always approachable with questions ranging from the most basic to the most off-the-wall (and he never held back from saying “that won’t work,” which was remarkably helpful). Luc was always there to bounce ideas off of, and provided constant support that my project was interesting, and would work, even when I wasn’t sure of either.

Of course, more day to day support came from my great labmates. Ben Lev introduced me to atom chips, and showed me the way around the clean room, which set me on the course for my initial project. I could not have designed and built that apparatus without patient answers and encouragement from John Stockton and Michael Armen. It was a pleasure to share an optical table with Kevin McHale, who also maintained the servers on which my code depends. Joe Kerckhoff welcomed me to his home on my trips to Stanford, making those trips much more pleasant and productive. Tony Miller and the rest of the lab made after-lunch foosball fun for years.

Once the lab started its move to Stanford, Oskar Painter and his students (particularly Paul, Raviv, Jessie, Chris, and Matt) welcomed me as an interloper into their space. While we did not end up getting things up and running together, I appreciated having an experimental home during the transition. In Berkeley, Birgitta Whaley and her group (Kevin, Chris, Raisa, and Mohan), graciously welcomed me into their space and group while I finished this research.

Sheri Stoll was a constant friendly presence in my years at Caltech, and always made sure the trains ran on time, travel reimbursements were processed, grant accounts were ready, and group events were scheduled and fed.

My parents, Gayle Muenchow and Richard Hopkins, provided years of priceless love and support. Mary has always been there for me. Thank you.

Abstract

Many quantum mechanical systems require large (potentially infinite) numbers of variables to exactly describe their state. In this thesis, I examine two approaches to develop simple, approximate models for such systems, which capture their essential dynamics. I use two bistable regimes of the Jaynes-Cummings model of cavity quantum electrodynamics as example systems to evaluate the effectiveness of each approach. In the phase bistable regime (which occurs with large driving field, and which I study in an on-resonance “bad cavity” regime to make numerical simulations tractable), the cavity field switches between two states with identical amplitude but opposite phase. In the absorptive bistable regime (which I study with small driving field in an on-resonance “good cavity”), two stable regions of state space differ in cavity field amplitude as well as their shape and qualitative behavior. After introducing these two regimes and their dynamics, I give a short introduction to projecting dynamical equations onto linear subspaces. Proper Orthogonal Decomposition (POD) allows the algorithmic construction of subspaces onto which the dynamics may be projected. I demonstrate that the application of POD to phase bistability results in effective approximate filters, while the asymmetry of the absorptive bistable case requires extensions to POD, developed in this thesis, to create a functional filter. Local Tangent Space Alignment is one of a class of unsupervised manifold learning algorithms which use the local geometry of high-dimensional data, such as quantum trajectories, to calculate the coordinates of that data on a low-dimensional manifold. I show how this algorithm functions, and characterize the manifolds that result from phase and absorptive bistability. I fit the 3-dimensional phase bistable manifold with a small set of system observables, and create a three-dimensional set of equations (similar to the semi-classical Maxwell-Bloch equations) which perform very well as a filter. Absorptive bistability again proves to be more complicated, but I am able to show that the underlying manifold is small, and make some progress on characterizing its relations with system observables.

Contents

Acknowledgments	iii
Abstract	v
Contents	vii
Preface	1
1 Introduction and Background	5
1.1 Thesis overview	7
1.2 Model system: Cavity quantum electrodynamics	8
1.2.1 The Jaynes-Cummings model	8
1.2.2 Filtering	10
1.3 The Maxwell-Bloch equations	11
1.4 Bistability and other interesting dynamics	13
1.5 An aside on stochastics	17
2 Projection Filters	19
2.1 Filter projection in general	19
2.2 Projecting onto a linear density matrix space	21
2.2.1 The stochastic master equation	21
2.2.2 The density matrix	22
2.2.3 Projection	23
2.2.4 The deterministic terms $A[\rho]$	24
2.2.5 The stochastic terms $B[\rho]$	27
2.2.6 Stratonovich back to Itô	29
3 Linear Models from Proper Orthogonal Decomposition	31
3.1 Background	31
3.2 Proper Orthogonal Decomposition of quantum trajectories	32

3.3	Phase bistability	35
3.3.1	Projection filters	36
3.4	Absorptive bistability	40
3.4.1	Weighted POD	43
3.4.2	Zoned POD	45
3.5	Discussion	50
4	Dynamics on Nonlinear Manifolds	53
4.1	Local Tangent Space Alignment	53
4.2	Deriving dynamical systems on LTSA manifolds	57
4.3	Phase bistability: Dynamics on three dimensional manifolds from LTSA	59
4.3.1	A new set of “Maxwell-Bloch” equations	63
4.4	Two and three dimensional manifolds for absorptive bistability	68
4.4.1	Fitting to observables	71
5	Conclusion	73
	Bibliography	77

Preface

The work contained in this thesis reflects only the final year of my graduate study at Caltech. Prior to this work, I spent five years as an experimentalist, working mostly on atom chips.

My graduate career began in the fall of 2002, when I arrived at Caltech after a year of post-baccalaureate research at Los Alamos National Laboratory. At Los Alamos, I worked with Salman Habib, who introduced me to the field of quantum optics, and quantum feedback control in particular. I worked with Salman, Kurt Jacobs (a Los Alamos postdoc), and Keith Schwab, then at the National Security Agency's Laboratory for Physical Sciences, to explore the possibilities for feedback cooling of nanomechanical resonators. During that year, I also had the opportunity to meet several members of MabuchiLab at conferences, and the intellectual and personal links that drew me to the group were forged. In my first year at Caltech, I informally joined Hideo's group, and worked on finishing and submitting the nanomechanical cooling paper [1].

Once the paper was safely submitted and on its way to publication, I started to look for a home for myself within our group. I quickly gravitated to atom chips, and the experiments undertaken by Ben Lev. Rather than apprentice myself directly to Ben, however, Hideo suggested that I take an experiment which Ben had recently performed and look for ways to extend it. Ben had built an atom mirror out of an old hard drive by etching away stripes of magnetic material and magnetizing the remaining bands. This forms a magnetic field which decays exponentially away from the surface, which repels atoms which are in a low-field-seeking state. With this result in the toolbox, I began to look for ways to use such a mirror to build more complex devices. I quickly became attracted to the idea of magneto-electrostatic devices, which use the attractive DC Stark effect to balance the repulsive magnetic force. After many explorations of various geometries, I settled on the disk geometry shown in Figure 1a. With a geometry designed, and its predicted atomic potential calculated, we wrote a paper proposing the trap [2], and I set about building both the chips and an apparatus with which to use them.

Ben introduced me to the cleanroom facilities of Prof. Michael Roukes, where I learned the basics of chip fabrication from Ben and others. I designed a photomask for the final chip design, and set about devising a recipe which would allow the creation of an atom mirror with a 3 micron grating period, circular gold disks of carefully-determined thickness, and leads to charge the disks.

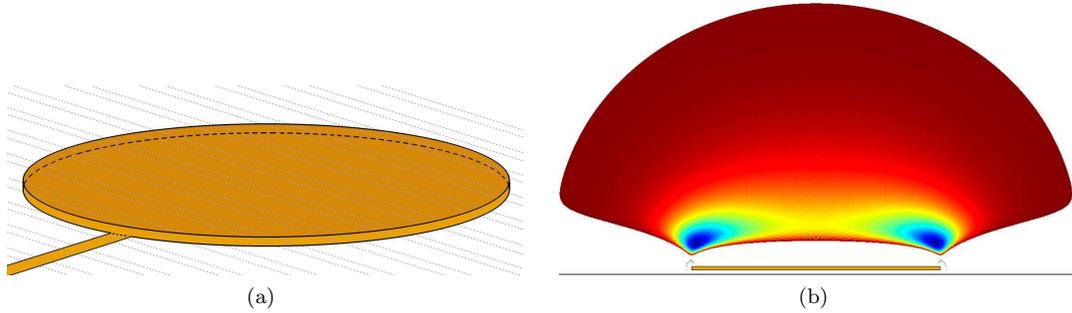


Figure 1: a) Schematic for a magneto-electrostatic ring trap. b) Trapping potential for a magneto-electrostatic trap, showing the sharp minima above the disk edge and the weaker trapping potential across the full disk surface.

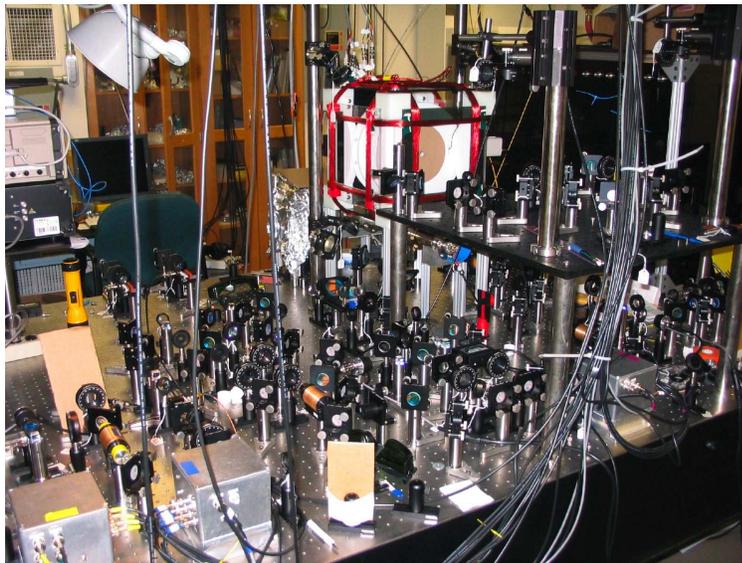


Figure 2: Rubidium cooling and trapping apparatus at Caltech.

At the same time, I claimed a section of optical table and began to design and build a laser cooling and trapping apparatus. I was the first in the lab to use rubidium atoms, rather than cesium, which generally made the experiment somewhat easier to build. First, the relevant light has a wavelength of approximately 780 nm, which is almost visible — much closer to visible than the 852 nm light used in cesium experiments. In addition, I was able to use getters to dispense rubidium into the vacuum system, rather than the cumbersome ampoules necessary for cesium. With the guidance of Ben Lev, John Stockton, and Michael Armen, and with computer input-output libraries built by JM Geremia, I was able to build the full apparatus, and cool and trap atoms into a magneto-optical trap (MOT) in March of 2006. A snapshot of the apparatus from around this time is shown in Figure 2.

After the publication of [2], I made several important changes in the way the experiment would be run, which in turn affected the apparatus I built. The critical change was in the trap loading. In [2], we proposed dropping in the atoms from a mirror MOT onto the atom chip. This would mean

dropping the atoms about a millimeter or two, which would increase our ability to load the traps, but would greatly restrict the number of atoms we could capture in the MOT cloud. I subsequently determined that we could load more atoms by inverting the chip and tossing the atoms up at it from a free-space MOT. The ability to capture up to 1,000 times more atoms in the initial MOT ($\sim 10^9$ rather than $\sim 10^6$) more than compensated for the reduction in loading efficiency by a factor of roughly 100. I simulated the launching of classical particles from a Gaussian MOT distribution and found that we could load different levels within the chip trap by changing the time at which we turned on the conservative trapping potential. Inverting the chip also made it easier to imagine how we would detect the trapped atoms: when the DC potential is removed, the atoms will fall away from the surface, where they can be imaged by fluorescence.

A significant downside of the inverted operation is that atoms would travel quite far (several inches) from the trapping MOT to the chip, and therefore needed to be very cold to maintain density. This required sub-Doppler cooling as a part of the atom-tossing procedure. After several months of system optimization (which made the whole apparatus more robust), I was finally able to cool the atoms to below their Doppler temperature one day in the spring of 2007. The next day, however, the magnetic field was distinctly different, and sub-Doppler cooling no longer worked. After scouring the lab for what had changed, I stepped next door into the Eisenstein lab, where I learned that their 13 Tesla superconducting magnet was at a different set-point that day than the day before: culprit identified. I was able to cobble together a sensor above the MOT to measure the background field from their magnet, but stopped work on optimizing that sensor when Hideo announced the lab move to Stanford. My apparatus, built using shared grant money with Prof. Oskar Painter, would stay at Caltech with me, but move to the Painter lab.

As my lab-mates prepared to move to Stanford, I disassembled and packed up my apparatus for the move across campus. I was initially supposed to reassemble in one lab, but there was not enough space. Then I designed the conversion of an office to a lab space, but that was not completed because Prof. Painter won space in a different building, for which he had to give up the space where I was planning to work. While this uncertainty plagued my progress, I continued the design process for an alternative experiment which would have a greater chance of success than the magneto-electrostatic traps (where the trap loading issue was still going to be very challenging to resolve): plasmonic diffraction gratings for atoms. Standing surface electromagnetic waves, at very short wavelengths due to their confinement in the surface, would allow atoms to be diffracted to large angles. However, as it became clear that I was not going to be able to reassemble my apparatus in time to get results and graduate in a reasonable time, I approached Hideo about a theoretical project, and this thesis project was born. The apparatus I designed and built has been reassembled in the Painter group's new space, but has not yet been fully integrated into their research.

The final year, almost to the day, of my graduate study was consumed by the theoretical and

computational work described in the following pages. My years as an experimentalist shaped my approach to these problems, always drawing my attention to how these calculations may be of use for physical experiments, but I also enjoyed the freedom from the many “messy” aspects of physical implementation.

Chapter 1

Introduction and Background

Quantum mechanics forms the basis for our understanding of the behavior of light and matter. While most day-to-day events we observe can be accurately modeled using classical mechanics and electromagnetism, many physical systems require a quantum mechanical description. Whether we are interested in the details of a chemical reaction or the flow of electrons through a semiconductor inside a computer, or even a system as apparently simple as the interaction of the electromagnetic field with a single atom, quantum mechanical modeling is both highly accurate and necessary. As optical networks and computing resources shrink and become faster, they will inevitably encounter limits set by quantum mechanics. While these limits might impose restrictions on the behavior of a system, they may also open the door to other uses for the same system. This thesis is dedicated to examining ways of building simpler models for quantum systems which respect their dynamics, but may eventually also allow those systems to be more easily developed into useful technology.

Quantum mechanics is an inherently probabilistic model for the physical world. In order to model the dynamics of quantum systems, we propagate not a single position and momentum as we might for a classical system, but instead a probability distribution. Undergraduate quantum mechanics usually introduces the Schrödinger equation, a partial differential equation for the wavefunction of an isolated physical system, from which we can calculate the possible results of measurements of certain observable quantities. However, as we work to match this simple model with the physical world, we encounter limits to the model. First, we are forced to reckon with the difficulty of building an isolated system, and then measuring it. By definition, the measurement process is an intrusion from the “outside” into a supposedly isolated system. This is usually dealt with by hand-waving arguments about an omnipotent experimenter suddenly introducing a measurement apparatus and making a sharp, projective measurement.

Of course, this model, too, often does not match with the physical reality of a laboratory or its example quantum system under study. A real physical quantum system is an open system, unavoidably coupled to the environment surrounding it, making it difficult to define the boundary of the “system.” Often, the best we can hope for is to shape the coupling of system to environment

so that we may retrieve information about the system without disrupting it in ways we do not want. (If we can do this, then we are justified in our model’s separation between “system” and “environment.”) A physical measurement process can almost always be modeled not as a direct measurement of the system, but instead a measurement of the environment interacting with the system of interest. Measurements also necessarily take time; the instantaneous measurement of introductory quantum mechanics is as much an approximation as its isolation.

We must also deal with the fact that our measurement may not tell us everything about our system, because of the limited way in which the environment interacts with the system, or because our measurement cannot distinguish between different states of the environment or, as a result, the system. As experimenters, our understanding of the state of our quantum system is shaped by the system’s history and dynamics (the realm of the Schrödinger equation), and also by the fallible way we measure the system, noise which may be introduced through the environment or measurement, or our simple failure to measure all of the parts of the environment which carry information about the system. We must almost always think of our quantum system as being in a mixed state, acknowledging our lack of knowledge about the system.

Mixed states cannot be modeled as wavevectors ($|\psi\rangle$), and must instead be modeled using density matrices ($\sum_i c_i |\psi_i\rangle\langle\psi_i|$), which we propagate in time using the master equation. Density matrices reside in a much larger space than wavefunctions — they have many more degrees of freedom — which makes accurate simulations of dynamics a challenge for large systems. In particular, any system which is coupled with an electromagnetic field can be difficult to simulate because the field, modeled as a harmonic oscillator, is infinite dimensional (there’s no top to the ladder of states). It is usually reasonable to define a cut-off energy, but this can still result in very large density matrices, especially if the system consists of tensor products of this large field space with other system components. Coupling with a two-level system quadruples the number of elements in the density matrix.

One way to tackle this challenge is by simulating only fully observed systems. That is, only systems in which the experimenter measures every output channel. Such systems can be simulated using a stochastic Schrödinger equation, and their states remain pure, and take the form of wavefunctions. The quantum simulations (“trajectories”) in this thesis are all of this form. The statistical behavior of such systems is identical to partially observed systems over long times (or multiple samples of identical systems). Measurements should not be able to change the fundamental statistics of the underlying system.

Given a physical system, however, the experimenter often does not have the option of measuring every output channel. An atom may emit light into all 4π steradians, for example, which cannot all be covered in detectors. To build her best guess as to the system state, the experimenter is forced to use the master equation as a “filter” — an equation which propagates her current best guess about the

state of the system, with filter input from the output of the measurement process. The experimenter’s best guess is a mixed state, and propagating the master equation will be computationally intense (and very difficult to do in real time). Should the experimenter desire to use this filter to decide on a control signal to drive the system to a particular state, she will be hard pressed to update her best guess in real time, and the control task will be very challenging. This thesis attempts to develop techniques to build simple models for quantum systems which capture the dynamics of interest (and potential utility) and which may also be propagated much more easily, potentially in real time, by a computer or purpose-built circuit. Such simple models might also provide insight into the system by elucidating the components which result in particular behavior.

The general technique I use to build simpler models is to find a linear or nonlinear submanifold of the full space of possible dynamics, in which the particular dynamics of a system are generally confined. I adapt various techniques for finding such subspaces developed for other applications to the case of quantum dynamics, and illustrate both their successes and failures.

1.1 Thesis overview

The remainder of this Introduction introduces the example physical system whose dynamics we will attempt to model with simple dynamical systems: a two-level atom in a high-finesse optical cavity, a situation known as “cavity quantum electrodynamics” (cavity QED). I review the dominant model for cavity QED, the Jaynes-Cummings model, give an abbreviated derivation of the Maxwell-Bloch equations, and a short summary of known interesting dynamics observed in these equations and in corresponding simulations of cavity QED. I close with a brief introduction to a few topics in stochastic calculus, to lay the groundwork for the remainder of the thesis.

Chapter 2 provides an overview of the technique of projecting filtering equations onto manifolds. I then turn to a particular manifold — a linear space of density matrices — and project the stochastic master equation for cavity QED onto this space, deriving nonlinear dynamical equations for the local coordinates.

Chapter 3 makes the work of Chapter 2 more concrete by describing a process, Proper Orthogonal Decomposition, for determining a linear density matrix space onto which to project the dynamics. I analyze cavity QED dynamics in phase and absorptive bistability regimes, and demonstrate some successes, and some failures, of this process for generating accurate filters.

Chapter 4 turns away from linear manifolds, and introduces nonlinear manifolds generated with an algorithm for Local Tangent Space Alignment. We turn again to phase and absorptive bistability regimes. In the phase bistability case, we are able to fit the manifold coordinates with simple combinations of expectation values, and derive a set of situation-specific “Maxwell-Bloch” equations, which I then compare with previously-derived equations for this system, and evaluate as a filter. The

absorptive bistability case is more complex, but I lay the groundwork for deriving similar systems of equations.

The concluding Chapter 5 draws the results together, examines the successes and failures of the examined model reduction techniques, and suggests directions for future work.

1.2 Model system: Cavity quantum electrodynamics

A single atom in a high-finesse optical cavity constitutes a canonical system in quantum optics. Experimental work on such systems (such as [3, 4, 5, 6, 7], among many others) has demonstrated strong coupling between the atom and the optical field, meaning that the presence or absence of a single photon drastically changes the environment for a single atom, and correspondingly, that the state of the single atom strongly affects the behavior of the field in the cavity. There are multiple parameter ranges in which strong coupling occurs. Strong coupling can be identified as a regime in which the ratio of the square of the atom-field coupling rate to the product of the cavity decay and atomic spontaneous emission rates is large. This ratio is called the “cooperativity.” Two limits of interest are the “bad cavity” and “good cavity” limits: for a fixed value of the cooperativity, the cavity decay rate may be large compared with the spontaneous emission (“bad cavity”), or small (“good cavity”). We usually scale time by the atomic spontaneous emission rate, so it ends up dropping out; this means that a small cavity decay rate, for a fixed cooperativity, implies a small atom-field coupling rate (inside a higher-finesse cavity), and *vice versa*.

1.2.1 The Jaynes-Cummings model

A simple model for the atom-cavity system is the Jaynes-Cummings model [8]. In this model, the atom is approximated as a two-level system, there is only one harmonic mode of the field in the cavity. The two-level atom is equivalent to a single spin, and the operators which act on it are the Pauli matrices and their linear combinations: σ_- lowers the atom into its ground state, while $\sigma_+ = \sigma_-^\dagger$ excites the atom. The field is acted upon by a and a^\dagger , the familiar annihilation and creation operators for a simple harmonic oscillator. We operate in a rotating frame, at the frequency of the driving field. The Hamiltonian in this model takes the form

$$H = \Delta_c a^\dagger a + \Delta_a \sigma_+ \sigma_- + ig_0 (a^\dagger \sigma_- - a \sigma_+) + i\mathcal{E} (a^\dagger - a), \quad (1.1)$$

where Δ_c is the detuning between the field and the cavity, Δ_a is the detuning between the field and the atomic transition, g_0 is the coupling rate between the atom and the cavity field, and \mathcal{E} is the driving (classical) field strength. Stepping through the terms of the equation:

- $\Delta_c a^\dagger a$ accounts for the difference in energy between on- and off-resonant drive of the cavity by the external field.
- $\Delta_a \sigma_+ \sigma_-$ similarly accounts for the difference in energy between on- and off-resonant drive of the atom by the external field.
- $ig_0 (a^\dagger \sigma_- - a \sigma_+)$ is responsible for exchange of energy between the atom and the field at rate $2g_0$. The first term creates an excitation in the field, while driving the atom to the ground state; that is, the atom emits a photon into the field. The second is the inverse process, where the atom absorbs one photon out of the field, and becomes excited.
- $i\mathcal{E} (a^\dagger - a)$ is the external driving field, acting to excite the phase quadrature of the cavity mode.

A Hamiltonian like (1.1) would suffice if we were concerned only with the unobserved dynamics of a closed quantum system. However, for useful systems, the “closed” model will not suffice. Instead, we must extend the picture to include the system’s interaction with its environment, and the behavior of an observer making measurements on the system. The observer may know the quantum state fully, or imperfectly, and so we model the dynamics of such an “open quantum system” with a master equation, which propagates the motion of a density matrix. The interaction of the system and its environment is governed by probabilities, so that the time at which the atom or field changes state (such as emitting a photon) is random. We therefore require a stochastic model for the propagation of the density matrix, which allows for the introduction of noise resulting from this inherently probabilistic behavior.

The Itô form of the stochastic master equation (SME) which governs the behavior of an atom-cavity system being observed with homodyne measurements of both the leaking cavity field and the atomic emission is

$$\begin{aligned}
d\rho = & -i[H, \rho]dt + \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) dt \\
& + \gamma (2\sigma_- \rho \sigma_+ - \sigma_+ \sigma_- \rho - \rho \sigma_+ \sigma_-) dt \\
& + i\sqrt{2\kappa} (\rho a^\dagger - a\rho - \text{Tr}[\rho (a^\dagger - a)]\rho) dW_1 \\
& + i\sqrt{2\gamma} (\rho \sigma_+ - \sigma_- \rho - i\text{Tr}[\rho (\sigma_+ - \sigma_-)]\rho) dW_2.
\end{aligned} \tag{1.2}$$

Here dW_1 and dW_2 are uncorrelated Wiener processes corresponding to noise on the two different measurement processes, κ is the rate of field decay out the end of the cavity (usually toward a detector), and γ is the rate of decay of the atomic state (atomic spontaneous emission) measured by a second homodyne detection process. γ_\perp , not included above, is the atomic transverse relaxation rate, and we scale time so that this rate is 1. A free atom has a spontaneous emission rate $\gamma = 2$, and

a Fabry-Perot cavity of the sort we model does not significantly change the spontaneous emission characteristics of the atom. The cooperativity parameter mentioned above is defined as $C = \frac{g_0^2}{2\kappa\gamma_\perp} = \frac{g_0^2}{\kappa\gamma}$.

This stochastic master equation will maintain the purity of an initial state, because all the channels of information leaving the system are being measured. As a result, this equation is exactly equivalent to a stochastic Schrödinger equation (SSE), which is relatively easy to simulate due to its more benign linear scaling behavior with the number of field modes to be propagated. I will refer to the pure states which result from the propagation of such an equation as *quantum trajectories*, and they will be a prime source of our insight into system behavior from simulations throughout this thesis. My discussion of quantum trajectories here is grounded in the work of Mabuchi and Wiseman [9] and references therein, although for convenience I have created a somewhat narrower definition of quantum trajectories as pure state trajectories only. I calculate these trajectories throughout the thesis using the SSE integration built in to the Quantum Optics Toolbox for Matlab, written by Sze Tan [10].

The experimentalist is generally unable to measure the field leaking out the sides of the cavity, and has only the measurement record from the cavity field measurement. As a result, she must average over all possible measurement results for the missing atomic measurement, leaving a (likely) mixed state as her best state of knowledge about the system. The density matrix in this case evolves according to the following SME:

$$\begin{aligned} d\rho = & -i[H, \rho]dt + \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) dt \\ & + \gamma (2\sigma_- \rho \sigma_+ - \sigma_+ \sigma_- \rho - \rho \sigma_+ \sigma_-) dt \\ & + i\sqrt{2\kappa} (\rho a^\dagger - a\rho - \text{Tr}[\rho (a^\dagger - a)]\rho) dW. \end{aligned} \quad (1.3)$$

It is this equation which we will project onto low-dimensional subspaces to produce reduced-order models (which we will then use as simple filters), and we will also use it to derive stochastic equations of motion for expectation values of system operators.

1.2.2 Filtering

Until now, we have simply discussed the stochastic master equation as a stochastic differential equation, without explicitly stating what we mean by the “state” ρ , or exactly how the noise on the observer’s measurement record relates to the Wiener process W driving the noise (dW) terms in the equation. We could, if we chose, continue while only thinking of this equation as the abstract dynamics of some theoretical system. However, it is far more useful to think of the stochastic master equations, Eqns. (1.2) and particularly (1.3), as filtering equations. That is, that the state ρ which we are propagating is the experimenter’s state of knowledge of the system. In this context, dW

represents the new information which the experimenter receives through her measurement, and is called the *innovation*. The innovation is defined as the difference between the measurement result and what the observer expected the result to be. In the case of a homodyne measurement of the phase quadrature of the cavity field, for example, the innovation is *measurement* – *expectation* = $dY - \frac{i}{2}\langle a^\dagger - a \rangle$. That this is a Wiener process reflects that the innovation is unbiased — the experimenter learns as much from a measurement result that is larger than her expectation as from a result that is smaller [11].

1.3 The Maxwell-Bloch equations

In my efforts to derive simple classical (or semi-classical) models to approximate quantum systems, the canonical relevant example of such a simplification serves as a constant point of comparison. The Maxwell-Bloch equations are a set of five, coupled, deterministic differential equations for the expectation values of the five simplest operators in this system: a , a^\dagger , σ_- , σ_+ , and σ_z . The first four will be familiar from the Hamiltonian; the last, σ_z , corresponds to the population difference between the atomic excited and ground states, and allows the system of equations to close without stretching any approximations to the breaking point. These operators may be combined to create the Hermitian operators

$$\begin{aligned} x &= \frac{1}{2}(a + a^\dagger) \\ y &= \frac{i}{2}(a^\dagger - a) \\ \sigma_x &= \sigma_+ + \sigma_- \\ \sigma_y &= i(\sigma_- - \sigma_+). \end{aligned} \tag{1.4}$$

These Hermitian operators correspond to possible observables of the atom-cavity system. x is the cavity field's amplitude quadrature; y is its phase. Using the spin analogy with a two-level atom, the Pauli matrix σ_x is the atomic component along the x axis, and similarly for σ_y and σ_z .

In the density matrix formalism, the expectation value of an operator \mathcal{O} is given by $\text{Tr}(\mathcal{O}\rho)$. Therefore, the equation of motion of that expectation value is $d\langle\mathcal{O}\rangle = \text{Tr}(\mathcal{O}(d\rho))$. If we use the definition of $d\rho$ given by Eqn. (1.3) to calculate the equations of motion of the five operators above,

we derive this set of equations:

$$\begin{aligned}
d\langle a \rangle &= -\kappa(1+i\Theta)\langle a \rangle dt + g_0\langle \sigma_- \rangle dt + \mathcal{E}dt + i\sqrt{2\kappa}(\langle a(a^\dagger - a) \rangle - \langle a^\dagger - a \rangle \langle a \rangle) dW \\
d\langle a^\dagger \rangle &= -\kappa(1-i\Theta)\langle a^\dagger \rangle dt + g_0\langle \sigma_+ \rangle dt + \mathcal{E}dt + i\sqrt{2\kappa}(\langle a^\dagger(a^\dagger - a) \rangle - \langle a^\dagger - a \rangle \langle a^\dagger \rangle) dW \\
d\langle \sigma_- \rangle &= -\gamma(1+i\Delta)\langle \sigma_- \rangle dt + g_0\langle a\sigma_z \rangle dt + i\sqrt{2\kappa}(\langle \sigma_-(a^\dagger - a) \rangle - \langle a^\dagger - a \rangle \langle \sigma_- \rangle) dW \\
d\langle \sigma_+ \rangle &= -\gamma(1-i\Delta)\langle \sigma_+ \rangle dt + g_0\langle \sigma_z a^\dagger \rangle dt + i\sqrt{2\kappa}(\langle \sigma_+(a^\dagger - a) \rangle - \langle a^\dagger - a \rangle \langle \sigma_+ \rangle) dW \\
d\langle \sigma_z \rangle &= -2\gamma(1+\langle \sigma_z \rangle) dt - 2g_0\langle a^\dagger \sigma_- + \sigma_+ a \rangle dt \\
&\quad + i\sqrt{2\kappa}(\langle \sigma_z(a^\dagger - a) \rangle - \langle a^\dagger - a \rangle \langle \sigma_z \rangle) dW.
\end{aligned} \tag{1.5}$$

This system of equations, however, is not closed: it includes the expectation values of products of operators. If we make the *ad hoc* approximation that these operator products can simply be factored ($\langle AB \rangle \rightarrow \langle A \rangle \langle B \rangle$), we are effectively neglecting correlations between observables (especially the correlations between the atom and the field), of the sort we might expect to see in quantum-limited behavior. However, this dramatically simplifies the system of equations. In particular, it eliminates all of the stochastic terms:

$$\begin{aligned}
d\langle a \rangle &= -\kappa(1+i\Theta)\langle a \rangle dt + g_0\langle \sigma_- \rangle dt + \mathcal{E}dt \\
d\langle a^\dagger \rangle &= -\kappa(1-i\Theta)\langle a^\dagger \rangle dt + g_0\langle \sigma_+ \rangle dt + \mathcal{E}dt \\
d\langle \sigma_- \rangle &= -\gamma(1+i\Delta)\langle \sigma_- \rangle dt + g_0\langle a \rangle \langle \sigma_z \rangle dt \\
d\langle \sigma_+ \rangle &= -\gamma(1-i\Delta)\langle \sigma_+ \rangle dt + g_0\langle \sigma_z \rangle \langle a^\dagger \rangle dt \\
d\langle \sigma_z \rangle &= -2\gamma(1+\langle \sigma_z \rangle) dt - 2g_0(\langle a^\dagger \rangle \langle \sigma_- \rangle + \langle \sigma_+ \rangle \langle a \rangle) dt.
\end{aligned} \tag{1.6}$$

This resulting set of equations is known as the *Maxwell-Bloch equations*, and has its origins in a separate, semi-classical derivation [12]. The Maxwell-Bloch equations are derived by extending the optical Bloch equations (which characterize the excitation and coherence of a two-level atomic medium) to include interaction with a resonant (or near-resonant) coherent optical field, for example in a laser.

Now that this is simply a closed set of equations for deterministic, classical observables, the full machinery of classical dynamical systems analysis can be brought to bear on understanding the different regimes of dynamics modeled by these equations. The form of the correspondence between the calculated dynamics of the Maxwell-Bloch equations and the behavior of quantum systems with equivalent parameters remains an open question, and provides a significant portion of the motivation for the research described in this thesis.

1.4 Bistability and other interesting dynamics

The quantum dynamical system defined by the master equations (1.2) and (1.3) is known to exhibit several interesting behaviors, some of which correspond with the behavior of the Maxwell-Bloch equations (1.6). Of particular interest for this thesis are two parameter regimes in which the dynamics of both quantum and classical systems are “bistable.” In classical, deterministic dynamics, a bistable system is one with two stable (or asymptotically stable) equilibria, and the system settles towards one or the other depending on initial conditions [13]. For the stochastic and quantum cases in this research, I have adapted this term, and assigned it a more phenomenological definition: a system is bistable if its dynamics show it to have two zones of phase space in which it is relatively stable. Noise may drive the system from one stable zone to the other (and back). I will commonly refer to the two zones as states, although they may not correspond to individual quantum states. The zones may roughly correspond to the stable points of a bistable deterministic system, or they may not. A more rigorous definition is likely possible using the techniques and language developed in stochastic dynamical systems theory (for background and foundations, see [14], [15] and [16]), but I have chosen this phenomenological definition for simplicity.

Bistability is useful in the engineering of practical devices, in particular for switching and binary memory. It is intimately related to hysteresis, in which a system prepared through two (or more) different time-varying processes settles into different stable states for the same set of system parameters. Static memory for computing makes use of the bistable behavior of magnetic domains — prepared with strong fields in one direction, they hold that state until actively switched. Most useful bistable systems are relatively noise-free, which contributes to their utility. As performance demands increase, however, devices must become smaller and faster, pushing them into the limit where their behavior is affected by thermal noise, and potentially quantum fluctuations. Current engineered devices are many orders of magnitude more energetic than the relevant quantum limits, but novel technologies for communication and computing (such as those developed for quantum computation and key distribution) may eventually approach it. Stabilizing the two “stable” states of a quantum bistable system in real time would require a model of the underlying quantum dynamics which can be computed alongside the system itself. This thesis is, in part, an attempt to examine potential tools which can be used to make such models, and evaluate their accuracy and utility.

In this thesis, I examine two distinct types of bistability in the cavity QED system: phase bistability and absorptive bistability. In the phase bistable regime, the field mode maintains a fixed amplitude, but switches between two states with opposite phase, with the switching events corresponding to atomic spontaneous emission events [9]. In the absorptive bistable regime, the field switches between a low-amplitude state, very near the vacuum state, and a high-amplitude state. For the parameter regime examined here, the upper state (or, perhaps more accurately,

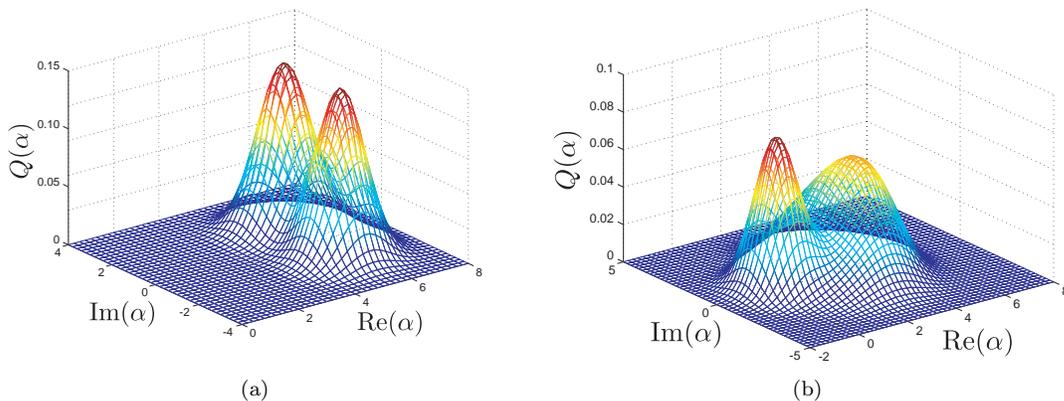


Figure 1.1: Q functions for the field modes of two bistable Cavity QED parameter regimes. (a) Phase bistability. (b) Absorptive Bistability.

zone) is much broader than a minimum-uncertainty state, and has significant variation in its phase. The appearance of the system from the observer’s standpoint, then, is a system whose absorption changes suddenly between two values, thus “absorptive bistability.” The relative ease of measuring field amplitude makes absorptive bistability particularly inviting for use as an optical switch or bit of memory. For both types of bistability, if one does a partial trace over the atomic portion of the density matrix, the Q function of the resulting field state takes the form of a bimodal distribution (Figure 1.1).

Phase bistability has been the subject of numerous publications over the course of the past twenty years. First noted by Alsing and Carmichael [17] and Kilin and Krinitskaya [18], it was further investigated by Mabuchi and Wiseman [9], who simulated the noisy switching behavior in a quantum trajectory formulation using the full stochastic master equation with homodyne detection to measure the phase quadrature directly, and no measurement of the spontaneous emission (following the example of Carmichael and collaborators [19], who modeled the system with homodyne detection of the cavity field and direct photo-detection of the atomic decay). Phase bistability occurs for resonant conditions (driving field, cavity, and atom all resonant), with a large driving field \mathcal{E} relative to the atom-cavity coupling g_0 , and g_0 large compared to the atomic spontaneous emission rate. In the limit of no atomic spontaneous emission ($\gamma = 0$), the Maxwell-Bloch equations have two stable points [17]

$$\begin{aligned}
 \langle a \rangle_{\pm} &= \frac{\mathcal{E} + g_0 \langle \sigma_- \rangle_{\pm}}{\kappa} \\
 \langle \sigma_- \rangle_{\pm} &= \frac{-g_0}{4\mathcal{E}} \mp i \left[\frac{1}{4} - \left(\frac{g_0}{4\mathcal{E}} \right)^2 \right]^{1/2} \\
 \langle \sigma_z \rangle &= 0,
 \end{aligned} \tag{1.7}$$

(with appropriate conjugate terms). Extending this insight into the case of quantum states, we see

that for $\mathcal{E} \gg g_0$, the two stable points become orthogonal quantum states. They now correspond to two coherent states of the field, each paired with an atomic state in a superposition of the excited and ground states, but with opposite phases. When we let $\gamma \neq 0$, but still small, the resulting spontaneous emission events correspond to the system switching between the two stable states [19].

Van Handel and Mabuchi [11] examined the phase bistable state from the perspective of projection filtering. They defined a three-dimensional, nonlinear manifold: one dimension measures the relative populations of two gaussian field states (paired with the corresponding atomic state appropriate for their phase), and the other two dimensions reflect the positions of the two gaussians. After an excellent, clear derivation of the filtering equation and rules for projecting it, they project the master equation onto this nonlinear manifold and derive a nonlinear set of stochastic differential equations to use as a simple filter. This filter behaves almost identically to the optimal filter, despite its simplicity, which implies that the underlying dynamics are fundamentally quite simple. If we take the two gaussian field modes to be fixed, the resulting 1-dimensional system is identical to the filter for a stationary Markovian jump process (the Wonham filter). Our analysis of the geometry of this system in Chapter 4 focuses on the geometry of the switching behavior; we should expect that the dynamics we derive will require the inclusion of the state position variables to characterize the transitions.

The Maxwell-Bloch equations may be examined as classical dynamical systems in order to search for regimes with interesting behavior which may correspond to novel behavior in the related quantum system. Gang, Ning, and Haken ([20] and [21]; see [21] for a summary of earlier related, but limited, work by others) undertook a search for these regimes, and Armen and Mabuchi [22] extended this search and analyzed the behavior of quantum systems in several regimes. They examined the absorptive bistability regime (which I use as an example system to evaluate model reduction techniques in this thesis), as well as behavior near both super- and sub-critical Hopf bifurcations, which lead to classical systems exhibiting limit cycles surrounding stable or unstable fixed points. Absorptive bistability may be predicted by semi-classical analysis of a saturable absorber in an optical cavity [23]. Savage and Carmichael [24] were the first to examine the single-atom case using the model described here, and discuss how quantum fluctuations would affect the system.

Absorptive bistability occurs for a range of parameters, generally with quite low driving field amplitude. Similar behavior with a non-resonant driving field is referred to as “dispersive bistability.” However, it is simplest to take all the detunings to be zero, and my calculations do so. The semiclassical system behavior then depends on the driving field strength and the dimensionless cooperativity. With $C = 10$, for example, the Maxwell-Bloch equations have two stable, and one unstable, solutions over a range of driving field strength; see Figure 1.2. The two stable states for a given driving field exhibit vastly different cavity field amplitude. Notably, the states are defined by the absolute value of their field amplitude, meaning that we do not gain much insight into the

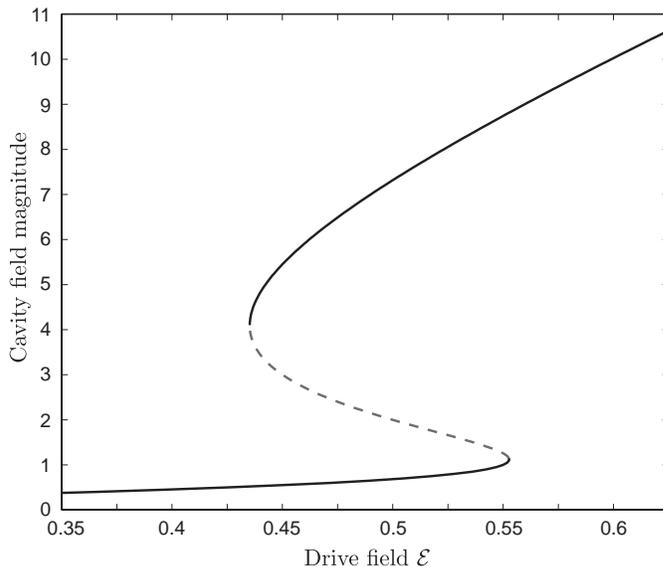


Figure 1.2: Semiclassical intracavity steady state field magnitude as a function of drive field \mathcal{E} , from [22]. The dashed portion indicates an unstable equilibrium. The parameter values are: $g_0 = \sqrt{2}$, $\gamma = 2$, $\kappa = 0.1$, $\Theta = 0$, and $\Delta = 0$ (with $\gamma_{\perp} = 1$ setting the scaling of time). The point used to examine absorptive bistability in this thesis has $\mathcal{E} = 0.56$.

relative quadrature (phase vs. amplitude) behavior. As [22] demonstrates, absorptive bistability parameters correspond to a quantum system which exhibits a bimodal Q function for the cavity field (see Figure 1.1b), and whose quantum trajectory simulations exhibit stochastic switching behavior.

In contrast to the phase bistable regime, however, these two “states” may not be not superficially identical. For the parameter regime of interest in this research, the lower state’s cavity field is very similar to a slightly displaced vacuum, while the upper state is much broader than a minimum-uncertainty coherent state, and may be better characterized as a semi-stable zone rather than a particular stable state. This shape is not universal. In the regime studied in [24], the two states are more similar in shape: the higher-amplitude portion of the Q function is only somewhat broader than the lower, and it lacks the spread in the phase quadrature that is pronounced with our parameters. For a fixed cooperativity, the shape of the Q function depends on the cavity decay rate, with more sharply peaked states in the good cavity case. Our trajectory simulations show a phase quadrature that stays very near zero in the lower state, and fluctuates dramatically in the upper state, corresponding to a pure state (as required by the quantum trajectory simulation) moving around significantly within the broad upper region defining the upper state. Making sense of this behavior, and what makes the upper state appear stable while not being a single state, motivates the examination of absorptive bistability throughout this thesis. While I find no global answer, I hope that the insights gained through attempts at model reduction will enable future work to better

understand this system.

1.5 An aside on stochastics

I will now give a brief overview of some results from stochastic calculus, to lay the groundwork for their use later in the thesis. These notes skim only the very surface of the field; I recommend Gardiner [25] and van Kampen [26] for more information. In stochastic calculus, we commonly have a Wiener process (usually denoted W), which is a continuous-time stochastic process which can be thought of as the integral of white noise. Rigorously, the differential form of writing stochastic equations of motion is nothing but a shorthand for the integral form. In stochastic differential equations, however, we write a stochastic increment dW , which is like a stochastic dt . We generally choose to scale it so that the mean value of dW^2 is dt , or “ $dW \sim \sqrt{dt}$.”

In classical, deterministic analysis, the integral can be defined as the sum of intervals, and in the limit of infinitesimal intervals, it does not matter whether we have chosen at each interval to use the function value at the start, middle, or end of the interval. In defining the stochastic integral, however, it does matter. We are forced to make a choice regarding which side of the infinitesimal interval to choose when summing, as they produce different results. The most common choices are those made by Itô and Stratonovich. Itô chose the start of each interval, while Stratonovich chose the midpoint. A given stochastic differential equation, therefore, must be accompanied by information about which kind of SDE it is. I follow the standard practice of writing $\circ dW$ for Stratonovich, and simply dW for Itô. Almost every equation in this thesis is in Itô form. There are many integrators designed for stochastic systems (see [27] for a summary), but I have restricted my work to a simple Itô-Euler integrator. (This is simply an Euler integrator with a stochastic term added.) When numerically integrating an SDE using this method, you must use the Itô form of the equation, because you are implicitly choosing the value at the start of each time interval for your integration. Other integrators use the Stratonovich form, or (like the Milstein integrator) add different correction terms.

Let us assume we have an Itô stochastic differential equation

$$dR(x, t) = A(x, t)dt + B(x, t)dW. \quad (1.8)$$

In order to transform this to the equivalent Stratonovich SDE, we must *subtract* the “Itô correction term”

$$\frac{1}{2} (\mathbf{D}B(x, t)) B(x, t)dt \quad (1.9)$$

where \mathbf{D} indicates the derivative. (If Eqn. (1.8) is a linear equation such as $d\mathbf{v} = A\mathbf{v}dt + B\mathbf{v}dW$, the derivative is simply the matrix B , so you may commonly see the correction term written with

B^2 .) Similarly, if we have a Stratonovich SDE

$$dR(x, t) = A(x, t)dt + B(x, t) \circ dW, \quad (1.10)$$

we *add* the Itô correction term (1.9) to transform it into Itô form.

The other critical difference between Itô and Stratonovich forms, for the purposes of this thesis, is how they respond to a transformation of coordinates. Projecting equations of motion, as I will do in Chapters 2 and 3, is such a transformation of coordinates. Under a transformation $\bar{x} = \phi(x)$, the coordinates of (1.8) transform into

$$\begin{aligned} \bar{A}(\bar{x}) &= A(x) \frac{d\phi}{dx} + \frac{1}{2} (B(x, t))^2 \frac{d^2\phi}{dx^2} \\ \bar{B}(\bar{x}) &= B(x) \frac{d\phi}{dx}. \end{aligned} \quad (1.11)$$

In contrast, the coordinates in the Stratonovich form transform like vectors (with no second-order derivatives):

$$\begin{aligned} \bar{A}(\bar{x}) &= A(x) \frac{d\phi}{dx} \\ \bar{B}(\bar{x}) &= B(x) \frac{d\phi}{dx}. \end{aligned} \quad (1.12)$$

Projection is a geometric process, so we need to be able to treat the stochastic increments in our equations of motion as vectors. Therefore, when we project equations of motion they will be Stratonovich equations [11].

Chapter 2

Projection Filters

Dynamics of open quantum systems take place in the space of density matrices, which can be a very high dimensional space, particularly when photon fields are involved. Strictly speaking, density matrices including photons are infinite, but it is common in practice to introduce a cutoff at some high Fock state, and work with finite, but large, density matrices. In simulations of cavity QED with the Jaynes-Cummings model, desired accuracy commonly requires one to keep track of 100 or more Fock states, in addition to the two-level atom. This results in a density matrix that is at least 200×200 , indicating that the dynamics take place in a space that is nominally 39,999-dimensional (an $N \times N$ density matrix has $N^2 - 1$ real degrees of freedom, taking Hermiticity into account). However, we know that the dynamics of a particular system do not fully explore this space, and as a result we would like to define a smaller (lower-dimensional) space, and study the system dynamics within that space alone. To do this, we project the equations of motion onto the lower-dimensional space. In this chapter I give a short overview of how to calculate these projected equations of motion, put this projection in the context of stochastic filtering equations, and then derive the form of the projected equations for the cavity QED master equation for a particular form of the lower-dimensional space: a linear density matrix space.

2.1 Filter projection in general

In this section, I will give a brief overview of the process of projecting equations of motion from a high-dimensional manifold onto a lower-dimensional one. I draw heavily upon the excellent description of this process given by van Handel and Mabuchi [11], adapting their derivation to the case of density matrices (instead of Q functions). Let us denote the space of all possible density matrices for a quantum system of interest by M , and the smaller subspace by S . Let our example stochastic dynamical system take the form

$$d\rho_t = A[\rho_t] dt + B[\rho_t] dW_t. \tag{2.1}$$

From a geometric standpoint, we would like to think of the right-hand side of this equation as a vector in the tangent space to M at a particular point θ_M . When we project the equation onto S , we would like to keep the components which are in the tangent space to S at θ , denoted $T_\theta S$, and discard the components in the orthogonal complement to this tangent space, denoted $T_\theta S^\perp$. For the right-hand side of Eqn. (2.1) to be treated as a vector, it must transform like one, which means we must interpret it as a Stratonovich stochastic differential equation, rather than an Itô equation. For now I will simply change the notation to reflect this, but in a practical situation (like that following, in Sec. 2.2), one would calculate the appropriate Itô correction term, which would change the forms of A and B . The new equation takes the form

$$d\rho_t = A[\rho_t] dt + B[\rho_t] \circ dW_t. \quad (2.2)$$

If we assume that we have a local coordinate system on S so that $\theta = (\theta_1, \theta_2, \dots)$, then we can write

$$T_\theta S = \text{Span} \left[\frac{\partial \rho(\theta)}{\partial \theta_1}, \frac{\partial \rho(\theta)}{\partial \theta_2}, \dots \right]. \quad (2.3)$$

We define an inner product on the space of density matrices

$$\langle \rho_A, \rho_B \rangle = \text{Tr}[\rho_A \rho_B], \quad (2.4)$$

which allows us to calculate the metric tensor in this basis:

$$\left\langle \frac{\partial \rho(\theta)}{\partial \theta^i}, \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle = \text{Tr} \left[\frac{\partial \rho(\theta)}{\partial \theta^i} \frac{\partial \rho(\theta)}{\partial \theta^j} \right] = g_{ij}(\theta). \quad (2.5)$$

If the basis defined in Eqn. (2.3) is orthonormal, g will simply be the Identity; otherwise it accounts for the non-orthonormality. With an inner product and a metric, we can define orthogonal projection of a vector field $X[\theta]$:

$$\Pi_\theta X[\theta] = \sum_i \sum_j g^{ij}(\theta) \left\langle X[\theta], \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle \frac{\partial \rho(\theta)}{\partial \theta^i}, \quad (2.6)$$

where g^{ij} denotes the (i, j) component of the inverse of the metric g defined in Eqn. (2.5).

We now wish to constrain the dynamics of Eqn. (2.2) to evolve on S :

$$d\rho(\theta_t) = \Pi_{\theta_t} A[\rho(\theta_t)] dt + \Pi_{\theta_t} B[\rho(\theta_t)] \circ dW_t, \quad (2.7)$$

which is just a stochastic differential equation for the parameters θ_t . Next, note that, in the Stratonovich calculus,

$$d\rho(\theta_t) = \sum_i \frac{\partial \rho(\theta_t)}{\partial \theta_t^i} \circ d\theta_t^i. \quad (2.8)$$

If we insert the definition of the orthogonal projection into Eqn. (2.2), we see that

$$\begin{aligned}
d\rho(\theta_t) &= \sum_i \sum_j g^{ij}(\theta) \left\langle A[\rho(\theta_t)], \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle \frac{\partial \rho(\theta)}{\partial \theta^i} dt \\
&\quad + \sum_i \sum_j g^{ij}(\theta) \left\langle B[\rho(\theta_t)], \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle \frac{\partial \rho(\theta)}{\partial \theta^i} \circ dW_t.
\end{aligned} \tag{2.9}$$

Comparing this expression with Eqn. (2.8), we can pull out the equations for $d\theta_t^i$:

$$d\theta_t^i = \sum_j g^{ij}(\theta) \left\langle A[\rho(\theta_t)], \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle dt + \sum_j g^{ij}(\theta) \left\langle B[\rho(\theta_t)], \frac{\partial \rho(\theta)}{\partial \theta^j} \right\rangle \circ dW_t. \tag{2.10}$$

Note that in order to apply this procedure, we need to know a functional form for $\rho(\theta)$, meaning we need a map from the smaller space (spanned by θ) to the larger space (where ρ lives). This is in addition to knowing the form of the projection from the larger space to the smaller, facilitated by Eqn. (2.6) and the like. (The manifold learning algorithms discussed in Chapter 4 provide only point-wise maps, so projecting the filters onto them will be a challenge.)

When we want to use these projected equations as a filter, the measurement photocurrent driving them is still the same as that which drives the full-space SME (thought of as a filter). The innovation process, dW , however, is different because it is defined as the difference between the measurement result and the filter's current estimate, which differs for each filter. Assuming that we can construct the map which reverses the projection Π , giving us a θ_M from each θ , we can directly compare the state generated by the projected equations of motion (2.10) with the corresponding trajectory. We should be careful to note that the projected equations will often be generated from an SME which does not correspond to measuring every output from the system, whereas a quantum trajectory simulation necessarily requires measurement of all outputs so as to allow the creation of a stochastic Schrödinger equation. In order to reduce the difference between these two cases for cavity QED, in trajectory simulations I have consciously chosen to measure the atomic spontaneous emission in the quadrature which gives the least additional information about the system in its measurement record (for both absorptive and phase bistability, this is the σ_y quadrature). It is possible that differences persist, but they ought to be minor because the trajectories are required to average (over long times or many runs) to the same mean as for the unmeasured-atom situation reflected in Eqn. (1.3).

2.2 Projecting onto a linear density matrix space

2.2.1 The stochastic master equation

The stochastic master equation we are concerned with is, as before, that for a two-level atom interacting with a single harmonic mode in an optical cavity, with measurement performed on the

field leaking out of the cavity. This is the normalized Itô form of the equation, for homodyne measurement of the phase quadrature:

$$\begin{aligned}
\mathcal{D}[\rho] = & -i[H, \rho]dt + \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) dt \\
& + \gamma (2\sigma\rho\sigma^\dagger - \sigma^\dagger\sigma\rho - \rho\sigma^\dagger\sigma) dt \\
& + \sqrt{2\kappa} (\rho a^\dagger - a\rho - i\text{Tr}[\rho(a^\dagger - a)]\rho) dW.
\end{aligned} \tag{2.11}$$

If we wanted to measure the amplitude quadrature instead, we would replace a with ia everywhere outside of the Hamiltonian. If we removed the nonlinear term $\text{Tr}[\rho(a^\dagger - a)]\rho$, we would have the un-normalized version of the SME. It has the distinct advantage of being linear, but will allow the trace of the density matrix to differ from 1. In a stochastic simulation, we can use the unnormalized equation, and simply renormalize ρ after each time step. However, for completeness, and because the filters from the normalized equation seem to be better “behaved,” I chose to use the normalized form, with its attendant complications resulting from nonlinearity.

In order for the geometry of projection to make sense, we need the components in this equation to transform like vectors, which means it needs to be a Stratonovich equation. We have two options for undertaking this transformation: 1) calculate the Itô correction term for Eqn. (2.11) or 2) use the much simpler (linear) un-normalized equation, transform it to Stratonovich form, and then normalize. I choose to do the first. This is the correct normalized Stratonovich form of the equation, calculated directly from Eqn. (2.11), for homodyne measurement of the phase quadrature:

$$\begin{aligned}
\mathcal{D}[\rho] = & -i[H, \rho]dt + \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) dt \\
& + \gamma (2\sigma\rho\sigma^\dagger - \sigma^\dagger\sigma\rho - \rho\sigma^\dagger\sigma) dt \\
& - \kappa \left(2a\rho a^\dagger - a^2\rho - \rho(a^\dagger)^2 + 2\text{Tr}[(a^\dagger - a)\rho](\rho a^\dagger - a\rho) \right. \\
& \left. - 2\rho (\text{Tr}[(a^\dagger - a)\rho])^2 + \rho [\text{Tr}[(a^\dagger - a)(\rho a^\dagger - a\rho)] \right] dt \\
& + i\sqrt{2\kappa} (\rho a^\dagger - a\rho - \text{Tr}[\rho(a^\dagger - a)]\rho) \circ dW,
\end{aligned} \tag{2.12}$$

where the Hamiltonian H is as in Eqn. (1.1), and dW is the innovation.

2.2.2 The density matrix

With a master equation in hand, we now turn to the possible forms of the space onto which we would like to project it. The dynamics of Eqn. (2.12) take place in the space of all density matrices (positive Hermitian operators with trace 1), but we expect that the dynamics of the system limit the fraction of this space which a physical system will explore. Proper Orthogonal Decomposition (see Chapter 3) generates a linear subspace directly from the dynamics, so we will now examine the

detailed form of such a space and the mechanics of the projection.

Imagine an N -dimensional linear density matrix space. Density matrices in this space have the following form:

$$\rho(v) = \rho_0 + \sum_{i=1}^N v_i \rho_i \quad (2.13)$$

where the ρ_i s are trace-0 Hermitian matrices (directions in density-matrix space), and ρ_0 is a positive, trace-1 Hermitian matrix (a valid density matrix), which serves as the origin in our linear space. The coefficients v_i are real, to maintain Hermiticity. There is nothing that forces $\rho(v)$ to remain positive, so it might cease to be a valid density matrix. However, when acting as part of a filter we expect it to stay positive almost all the time, except when presented with a measurement record which it is unable to do a good job of accommodating.

The partial derivatives of ρ are

$$\frac{\partial \rho(v)}{\partial v^i} = \rho_i. \quad (2.14)$$

We recall the definition of the inner product between matrices/operators as the trace of the product, and so we define the metric in this space

$$g_{ij} = \text{Tr}[\rho_i \rho_j] \quad (i, j > 0). \quad (2.15)$$

We assume that the ρ_i s have been orthonormalized so that $g_{ij} = g^{ij} = \delta_{ij}$ ($g = \text{Id}$).

If we had not used the normalized form of the stochastic master equation, we would have extended the dimension of the linear space by 1 to include a coefficient on ρ_0 . Then we would redefine the state to be the ratio of each coefficient to v_0 , which would complicate the equations to be evolved. Alternatively, in simulations, we would simply rescale all of the coefficients at each time step, setting $v_i = \tilde{v}_i / \tilde{v}_0$, $i \geq 0$. In practice, filtering using the normalized equations seems to be somewhat more robust, and it has the advantage of providing us with exact, nonlinear equations directly.

2.2.3 Projection

Following the general derivation given in Section 2.1, and specializing to our particular space S , spanned by the states ρ_i , we have that the orthogonal projection of (2.12) is

$$\Pi_v \mathcal{D}[\rho(v)] = \sum_{i=1}^N \sum_{j=1}^N g^{ij} \left\langle \mathcal{D}[\rho(v)], \frac{\partial \rho(v)}{\partial v^j} \right\rangle \frac{\partial \rho(v)}{\partial v^i}. \quad (2.16)$$

Simplifying because we know that $g_{ij} = g^{ij} = \delta_{ij}$, we see that

$$\Pi_v \mathcal{D}[\rho(v)] = \sum_{i=1}^N \langle \mathcal{D}[\rho(v)], \rho_i \rangle \rho_i. \quad (2.17)$$

For an arbitrary filtering SME of the form (2.2), we constrain the filter to evolve in our space of density matrices, and combine Eqn. (2.17) with Eqn. (2.7) to find that

$$dv_i = \langle A[\rho(v)], \rho_i \rangle dt + \langle B[\rho(v)], \rho_i \rangle \circ dW. \quad (2.18)$$

We have split the master equation, Eqn. (2.12), into deterministic and stochastic parts, as in Eqn. (2.2), to clarify calculation:

$$\begin{aligned} A[\rho] &= -i[H, \rho] + \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) \\ &\quad + \gamma (2\sigma\rho\sigma^\dagger - \sigma^\dagger\sigma\rho - \rho\sigma^\dagger\sigma) \\ &\quad - \kappa \left(2a\rho a^\dagger - a^2\rho - \rho(a^\dagger)^2 + 2\text{Tr}[(a^\dagger - a)\rho](\rho a^\dagger - a\rho) \right. \\ &\quad \left. - 2\rho(\text{Tr}[(a^\dagger - a)\rho])^2 + \rho[\text{Tr}[(a^\dagger - a)(\rho a^\dagger - a\rho)]] \right) \end{aligned} \quad (2.19)$$

$$B[\rho] = i\sqrt{2\kappa} (\rho a^\dagger - a\rho - \text{Tr}[\rho(a^\dagger - a)]\rho). \quad (2.20)$$

We will now project each of these terms, in order to derive the detailed form of Eqn. (2.18).

2.2.4 The deterministic terms $\mathbf{A}[\rho]$

Let us start calculating the terms in (2.18). First, we define

$$\mathcal{L}_H(\rho) \equiv -i[H, \rho] \quad (2.21)$$

$$\mathcal{L}_a(\rho) \equiv \kappa (2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) \quad (2.22)$$

$$\mathcal{L}_\sigma(\rho) \equiv \gamma (2\sigma\rho\sigma^\dagger - \sigma^\dagger\sigma\rho - \rho\sigma^\dagger\sigma) \quad (2.23)$$

$$\mathcal{L}_{ISL}(\rho) \equiv \kappa \left(2a\rho a^\dagger - a^2\rho - \rho(a^\dagger)^2 \right) \quad (2.24)$$

$$\begin{aligned} \mathcal{L}_{ISN}(\rho) &\equiv \kappa \left(2\text{Tr}[(a^\dagger - a)\rho](\rho a^\dagger - a\rho) - 2\rho(\text{Tr}[(a^\dagger - a)\rho])^2 \right. \\ &\quad \left. + \rho[\text{Tr}[(a^\dagger - a)(\rho a^\dagger - a\rho)]] \right). \end{aligned} \quad (2.25)$$

Then

$$\begin{aligned} \langle A[\rho(v)], \rho_i \rangle &= \langle \mathcal{L}_H(\rho(v)) + \mathcal{L}_a(\rho(v)) + \mathcal{L}_\sigma(\rho(v)) \\ &\quad - \mathcal{L}_{ISL}(\rho(v)) - \mathcal{L}_{ISN}(\rho(v)), \rho_i \rangle \end{aligned} \quad (2.26)$$

$$\begin{aligned} &= \langle \mathcal{L}_H(\rho(v)), \rho_i \rangle + \langle \mathcal{L}_a(\rho(v)), \rho_i \rangle + \langle \mathcal{L}_\sigma(\rho(v)), \rho_i \rangle \\ &\quad - \langle \mathcal{L}_{ISL}(\rho(v)), \rho_i \rangle - \langle \mathcal{L}_{ISN}(\rho(v)), \rho_i \rangle. \end{aligned} \quad (2.27)$$

Expanding the form of $\rho(v)$, we have

$$\langle \mathcal{L}_H(\rho(v)), \rho_i \rangle = \left\langle \mathcal{L}_H(\rho_0 + \sum_{j=1}^N v_j \rho_j), \rho_i \right\rangle \quad (2.28)$$

$$= \langle \mathcal{L}_H(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{L}_H(\rho_j), \rho_i \rangle, \quad (2.29)$$

and the same for \mathcal{L}_a , \mathcal{L}_σ , and \mathcal{L}_{ISL} , because they are all linear in ρ .

In fact, if we define

$$\mathcal{L} \equiv \mathcal{L}_H + \mathcal{L}_a + \mathcal{L}_\sigma - \mathcal{L}_{ISL} \quad (2.30)$$

then we have

$$\langle \mathcal{L}(\rho(v)), \rho_i \rangle = \langle \mathcal{L}(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{L}(\rho_j), \rho_i \rangle. \quad (2.31)$$

Plugging this into Eqn. (2.18), we see that the linear, deterministic part of dv is

$$dv_{i(\text{indet})} = \langle \mathcal{L}(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{L}(\rho_j), \rho_i \rangle dt. \quad (2.32)$$

If we think of dv as a vector, we see that this is just a matrix multiplication, where each entry in the matrix L is simply

$$L_{ij} = \langle \mathcal{L}(\rho_j), \rho_i \rangle + \langle \mathcal{L}(\rho_0), \rho_i \rangle \delta_{ij}. \quad (2.33)$$

Now we need to take a look at \mathcal{L}_{ISN} , the nonlinear terms from the Itô to Stratonovich conversion.

$$\begin{aligned} \mathcal{L}_{ISN}(\rho) \equiv & \kappa \left(2\text{Tr}[(a^\dagger - a)\rho](\rho a^\dagger - a\rho) - 2\rho (\text{Tr}[(a^\dagger - a)\rho])^2 \right. \\ & \left. + \rho [\text{Tr}[(a^\dagger - a)(\rho a^\dagger - a\rho)]] \right). \end{aligned} \quad (2.34)$$

Let us start with the first term, and plug in the approximate form of ρ for our linear space.

$$\begin{aligned} & \text{Tr} \left[(a^\dagger - a) \left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) \right] \left(\left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) a^\dagger - a \left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) \right) \\ = & \text{Tr}[(a^\dagger - a)\rho_0](\rho_0 a^\dagger - a\rho_0) + \text{Tr}[(a^\dagger - a)(\rho_0)] \left(\left(\sum_{j=1}^N v_j \rho_j \right) a^\dagger - a \left(\sum_{j=1}^N v_j \rho_j \right) \right) \\ & + \text{Tr} \left[(a^\dagger - a) \left(\sum_{j=1}^N v_j \rho_j \right) \right] ((\rho_0) a^\dagger - a(\rho_0)) \\ & + \text{Tr} \left[(a^\dagger - a) \left(\sum_{j=1}^N v_j \rho_j \right) \right] \left(\left(\sum_{j=1}^N v_j \rho_j \right) a^\dagger - a \left(\sum_{j=1}^N v_j \rho_j \right) \right) \end{aligned}$$

$$\begin{aligned}
&= \text{Tr}[(a^\dagger - a)\rho_0](\rho_0 a^\dagger - a\rho_0) + \text{Tr}[(a^\dagger - a)\rho_0] \sum_{k=1}^N v_k (\rho_k a^\dagger - a\rho_k) \\
&\quad + \sum_{j=1}^N v_j \text{Tr}[(a^\dagger - a)\rho_j](\rho_0 a^\dagger - a\rho_0) + \sum_{j=1}^N v_j \text{Tr}[(a^\dagger - a)\rho_j] \sum_{k=1}^N v_k (a\rho_k + \rho_k a^\dagger) \\
&= \left(\text{Tr}[(a^\dagger - a)\rho_0] + \sum_{j=1}^N v_j \text{Tr}[(a^\dagger - a)\rho_j] \right) \times \left((\rho_0 a^\dagger - a\rho_0) + \sum_{k=1}^N v_k (\rho_k a^\dagger - a\rho_k) \right). \tag{2.35}
\end{aligned}$$

Now, let us do the inner product with ρ_i , noting that the traces are things we have calculated anyway, because they're just the expectation values of $-2iy$ for each ρ_j :

$$\begin{aligned}
&\left(\text{Tr}[(a^\dagger - a)\rho_0] + \sum_{j=1}^N v_j \text{Tr}[(a^\dagger - a)\rho_j] \right) \\
&\quad \times \left\langle (\rho_0 a^\dagger - a\rho_0) + \sum_{k=1}^N v_k (\rho_k a^\dagger - a\rho_k), \rho_i \right\rangle \\
&= -2i \left(\langle y_0 \rangle + \sum_{j=1}^N v_j \langle y_j \rangle \right) \times \left(\langle (\rho_0 a^\dagger - a\rho_0), \rho_i \rangle + \sum_{k=1}^N v_k \langle (\rho_k a^\dagger - a\rho_k), \rho_i \rangle \right). \tag{2.36}
\end{aligned}$$

It doesn't simplify much because we can't use the orthogonality of the ρ_i s once the a s are present. For simulations, however, we can pre-calculate the values of everything in the angle brackets.

Let us take the third term, the other quadratic term:

$$\begin{aligned}
&\left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) \text{Tr} \left[(a^\dagger - a) \left(\left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) a^\dagger - a \left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) \right) \right] \\
&= \rho_0 \text{Tr}[(a^\dagger - a)(\rho_0 a^\dagger - a\rho_0)] + \rho_0 \sum_{k=1}^N v_k [\text{Tr}[(a^\dagger - a)(\rho_k a^\dagger - a\rho_k)]] \\
&\quad + \sum_{j=1}^N v_j \rho_j [\text{Tr}[(a^\dagger - a)(\rho_0 a^\dagger - a\rho_0)]] \\
&\quad + \sum_{j=1}^N v_j \rho_j \sum_{k=1}^N v_k [\text{Tr}[(a^\dagger - a)(\rho_k a^\dagger - a\rho_k)]] \\
&= \left(\text{Tr}[(a^\dagger - a)(\rho_0 a^\dagger - a\rho_0)] + \sum_{k=1}^N v_k [\text{Tr}[(a^\dagger - a)(\rho_k a^\dagger - a\rho_k)]] \right) \\
&\quad \times \left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right). \tag{2.37}
\end{aligned}$$

This does simplify once we do the inner product with ρ_i :

$$\langle \rho_0, \rho_i \rangle + \sum_{j=1}^N v_j \langle \rho_j, \rho_i \rangle = G_i + v_i \quad (2.38)$$

where

$$G_i \equiv \langle \rho_0, \rho_i \rangle. \quad (2.39)$$

And now for the cubic term:

$$\left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) \left(\text{Tr} \left[(a^\dagger - a) \left(\rho_0 + \sum_{k=1}^N v_k \rho_k \right) \right] \right)^2. \quad (2.40)$$

Noting the above simplification, we see that after we do the inner product, we have this:

$$\begin{aligned} & (G_i + v_i) \left(\text{Tr} [(a^\dagger - a)\rho_0] \right)^2 + 2 \sum_{k=1}^N v_k \text{Tr} [(a^\dagger - a)\rho_0] \text{Tr} [(a^\dagger - a)\rho_k] \\ & + \sum_{j=1}^N \sum_{k=1}^N v_j v_k \text{Tr} [(a^\dagger - a)\rho_j] \text{Tr} [(a^\dagger - a)\rho_k] \\ = & -4(G_i + v_i) \left(\langle y_0 \rangle^2 + 2 \sum_{k=1}^N v_k \langle y_0 \rangle \langle y_k \rangle + \sum_{j=1}^N \sum_{k=1}^N v_j v_k \langle y_j \rangle \langle y_k \rangle \right). \end{aligned} \quad (2.41)$$

Assembling all of the parts of $\mathcal{L}_{ISN}(\rho)$ together, we have:

$$\begin{aligned} dv_{i(ISN)} &= -\langle \mathcal{L}_{ISN}(\rho), \rho_i \rangle dt \\ &= -\kappa \left(-4i \left(\langle y_0 \rangle + \sum_{j=1}^N v_j \langle y_j \rangle \right) \times \left(\langle (\rho_0 a^\dagger - a \rho_0), \rho_i \rangle + \sum_{k=1}^N v_k \langle (\rho_k a^\dagger - a \rho_k), \rho_i \rangle \right) \right. \\ & \quad + 8(G_i + v_i) \left(\langle y_0 \rangle + \sum_{k=1}^N v_k \langle y_k \rangle \right)^2 \\ & \quad \left. + (G_i + v_i) \left(\text{Tr}[(a^\dagger - a)(\rho_0 a^\dagger - a \rho_0)] + \sum_{k=1}^N v_k [\text{Tr}[(a^\dagger - a)(\rho_k a^\dagger - a \rho_k)]] \right) \right) dt. \end{aligned} \quad (2.42)$$

2.2.5 The stochastic terms $B[\rho]$

Recall that

$$B[\rho] = i\sqrt{2\kappa} (\rho a^\dagger - a \rho - \text{Tr}[\rho (a^\dagger - a)] \rho). \quad (2.43)$$

The linear portion of this [$\mathcal{LS}(\rho) \equiv i\sqrt{2\kappa}(\rho a^\dagger - a\rho)$] is just like the deterministic case, with

$$dv_{i(\text{lin})} = \langle \mathcal{LS}(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{LS}(\rho_j), \rho_i \rangle \circ dW, \quad (2.44)$$

and will work out to a simple, constant over time, matrix multiplication by a matrix LS :

$$LS_{ij} = i\sqrt{2\kappa} (\langle \rho_j a^\dagger - a\rho_j, \rho_i \rangle + \langle \rho_0 a^\dagger - a\rho_0, \rho_i \rangle \delta_{ij}), \quad (2.45)$$

Now let us turn to the nonlinear term:

$$i \langle \text{Tr}[\rho(v)(a^\dagger - a)] \rho(v), \rho_i \rangle = i \left\langle \text{Tr} \left[\left(\rho_0 + \sum_{j=1}^N v_j \rho_j \right) (a^\dagger - a) \right] \left(\rho_0 + \sum_{k=1}^N v_k \rho_k \right), \rho_i \right\rangle. \quad (2.46)$$

This breaks out into 4 chunks: A constant (independent of v_i),

$$i \langle \text{Tr}[\rho_0 (a^\dagger - a)] \rho_0, \rho_i \rangle, \quad (2.47)$$

two linear terms:

$$i \left\langle \text{Tr}[\rho_0 (a^\dagger - a)] \left(\sum_{k=1}^N v_k \rho_k \right), \rho_i \right\rangle \text{ and} \quad (2.48)$$

$$i \left\langle \text{Tr} \left[\left(\sum_{j=1}^N v_j \rho_j \right) (a^\dagger - a) \right] \rho_0, \rho_i \right\rangle, \quad (2.49)$$

and one quadratic term:

$$i \left\langle \text{Tr} \left[\left(\sum_{j=1}^N v_j \rho_j \right) (a^\dagger - a) \right] \left(\sum_{k=1}^N v_k \rho_k \right), \rho_i \right\rangle. \quad (2.50)$$

Many of the components of these terms are constants.

The constant term is

$$2G_i \langle y_0 \rangle, \quad (2.51)$$

and the two linear terms are

$$2 \langle y_0 \rangle \sum_{k=1}^N v_k \langle \rho_k, \rho_i \rangle = 2 \langle y_0 \rangle \sum_{k=1}^N v_k \delta_{ik} = 2 \langle y_0 \rangle v_i \quad (2.52)$$

and

$$2G_i \sum_{j=1}^N v_j \langle y_j \rangle. \quad (2.53)$$

The quadratic term is

$$2 \sum_{j=1}^N \sum_{k=1}^N v_k v_j \langle y_j \rangle \langle \rho_k, \rho_i \rangle = 2 \sum_{j=1}^N \sum_{k=1}^N v_k v_j \langle y_j \rangle \delta_{ki} = 2v_i \sum_{j=1}^N v_j \langle y_j \rangle. \quad (2.54)$$

The stochastic part of Eqn. (2.18) from the nonlinear (trace) term is therefore

$$\begin{aligned} dv_{i(\text{trace})} &= -\sqrt{8\kappa} \left(\langle y_0 \rangle G_i + \langle y_0 \rangle v_i + G_i \sum_{j=1}^N v_j \langle y_j \rangle + v_i \sum_{j=1}^N v_j \langle y_j \rangle \right) \circ dW \\ &= -\sqrt{8\kappa} \left(\langle y_0 \rangle G_i + \langle y_0 \rangle v_i + (G_i + v_i) \sum_{j=1}^N v_j \langle y_j \rangle \right) \circ dW. \end{aligned} \quad (2.55)$$

So, putting all of Eqn. (2.18) together we have

$$\begin{aligned} dv_i &= \left(\langle \mathcal{L}(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{L}(\rho_j), \rho_i \rangle \right) dt + dv_{i(ISO)} \\ &\quad + \left(\langle \mathcal{LS}(\rho_0), \rho_i \rangle + \sum_{j=1}^N v_j \langle \mathcal{LS}(\rho_j), \rho_i \rangle \right) \circ dW \\ &\quad - \sqrt{8\kappa} (G_i + v_i) \left(\langle y_0 \rangle + \sum_{j=1}^N v_j \langle y_j \rangle \right) \circ dW. \end{aligned} \quad (2.56)$$

2.2.6 Stratonovich back to Itô

For numeric simulation with an Itô-Euler integrator, we require Itô equations, so we need to transform our Stratonovich equations back into Itô. Currently our equations of motion for the projected filter have the form

$$\mathcal{D}[\mathbf{v}] = A_v[\mathbf{v}]dt + B_v[\mathbf{v}] \circ dW. \quad (2.57)$$

The correction term has the form

$$\frac{1}{2} (\mathbf{D}B_v[\mathbf{v}]) B_v[\mathbf{v}] \quad (2.58)$$

where $\mathbf{D}(\cdot)$ is the derivative.

The \mathcal{LS} part of the stochastic term is just matrix multiplication by LS , so its derivative is just LS .

The part of B that comes from the normalizing term in the SME has a derivative of

$$\begin{aligned}
\mathbf{D}B_{v(\text{trace})}[\mathbf{v}]_{ij} &= -\frac{\partial}{\partial v_j} \sqrt{8\kappa} \left(\langle y_0 \rangle G_i + \langle y_0 \rangle v_i + (G_i + v_i) \sum_{k=1}^N v_k \langle y_k \rangle \right) \\
&= -\sqrt{8\kappa} \left(\langle y_0 \rangle \delta_{ij} + (G_i + v_i) \sum_{k=1}^N \langle y_k \rangle \delta_{jk} + \delta_{ij} \sum_{k=1}^N v_k \langle y_k \rangle \right) \\
&= -\sqrt{8\kappa} \left(\left(\langle y_0 \rangle + \sum_{k=1}^N v_k \langle y_k \rangle \right) \delta_{ij} + (G_i + v_i) \langle y_j \rangle \right). \tag{2.59}
\end{aligned}$$

For notational simplicity, let us call this matrix DB . Let us call the part of B that comes from the normalizing term, which is a vector, \mathbf{B}_n (you can read its elements off of Eqn. (2.55)). Then the full correction term to take us back to Itô form is

$$\frac{1}{2} (\mathbf{D}B_v[\mathbf{v}]) B_v[\mathbf{v}] = \frac{1}{2} (LS + DB) (\mathbf{L}\mathbf{S} + LS\mathbf{v} + \mathbf{B}_n) \tag{2.60}$$

where

$$\mathbf{L}\mathbf{S}_i = \langle \mathcal{L}\mathcal{S}(\rho_0), \rho_i \rangle. \tag{2.61}$$

Applying this term, we now have the complete Itô stochastic differential equation for the dynamics of the projected filter:

$$\begin{aligned}
d\mathbf{v} &= \left(\mathbf{L} + L\mathbf{v} + d\mathbf{v}_{ISN} + \frac{1}{2} (LS + DB) (\mathbf{L}\mathbf{S} + LS\mathbf{v} + \mathbf{B}_n) \right) dt \\
&\quad + (\mathbf{L}\mathbf{S} + LS\mathbf{v} + \mathbf{B}_n) dW \tag{2.62}
\end{aligned}$$

where

$$\mathbf{L}_i = \langle \mathcal{L}(\rho_0), \rho_i \rangle. \tag{2.63}$$

With the machinery in place, we can now turn to generating linear density matrix spaces by Proper Orthogonal Decomposition of quantum trajectories, project the filter onto them, and evaluate their performance.

Chapter 3

Linear Models from Proper Orthogonal Decomposition

Proper Orthogonal Decomposition (POD), alternatively known as Principal Component Analysis or the Karhunen-Loève decomposition, is a model-reduction technique which generates the optimal linear subspace of dimension D for a given set of higher-dimensional data. That is, if the data are contained within an attractor, the POD process can produce the affine linear space that best approximates the space containing that attractor. In this Chapter, I give a short derivation of the POD algorithm, then show the results from applying it to two different bistable atom-cavity regimes. For each regime, I show the performance of filters based on these POD results.

3.1 Background

The Proper Orthogonal Decomposition (POD) process has been derived in a variety of fields, which has resulted in it having several names. In fluid dynamics, it is known as the Karhunen-Loève decomposition, and has its origins in the work of Lumley [28]. The method itself had its origins in Pearson [29] and again in Hotelling [30].

The derivation which follows is based on those presented by Lall *et al.* [31] and Zhang and Zha [32]. We start with an empirical set of data — an unordered collection of N points $x^{(i)}$ in \mathbb{R}^m . In the cases to be examined below, these will be vectorized forms of the density matrix ρ . We would like to build the optimal d -dimensional affine subspace of \mathbb{R}^m , which will be isomorphic to \mathbb{R}^d . Zhang and Zha take the approach of building the optimal map from the lower dimensional space to the higher, while Lall *et al.* choose to find the optimal projector from the larger space to the smaller. Here I choose the latter path, because it is more consistent with our goal of projecting the dynamics into the lower-dimensional space.

We would like to find the optimal projection matrix, $Q \in \mathbb{R}^{m \times m}$, from \mathbb{R}^m to a d -dimensional subspace of \mathbb{R}^m . We define the optimal Q as that which minimizes the total squared perpendicular

distance of the original set of points from the d -dimensional plane:

$$E(Q) = \sum_{i=1}^N \|x^{(i)} - Qx^{(i)}\|^2. \quad (3.1)$$

We wish to find the optimal *affine* subspace, which must go through the mean of the data, so we subtract the mean of all the $x^{(i)}$ s, \bar{x} , from each point before proceeding. We next construct the correlation matrix

$$R = \sum_{i=1}^N (x^{(i)} - \bar{x})(x^{(i)} - \bar{x})^*, \quad (3.2)$$

and calculate its eigenvalues, ordered $\lambda_1 \geq \lambda_2 \geq \dots \lambda_m$. It is then a standard result that

$$\min_Q E(Q) = \sum_{l=m-d+1}^m \lambda_l. \quad (3.3)$$

The magnitude of each eigenvalue measures the relative contribution of the direction corresponding to the paired eigenvector to the data distribution as a whole. By cutting off the space at d dimensions, we measure the error, E , by summing the eigenvalues corresponding to the directions we have chosen to discard.

Turning now to constructing the projector into this optimal subspace, we make use of the orthonormal eigenvectors $\phi_1, \phi_2, \dots \phi_d$ corresponding to the largest eigenvalues. The approximate $\hat{x}^{(i)}$ to $x^{(i)}$ is given by

$$\hat{x}^{(i)} = \sum_{j=1}^d a_{ij} \phi_j + \bar{x} \quad (3.4)$$

where

$$a_{ij} = \langle x^{(i)} - \bar{x}, \phi_j \rangle. \quad (3.5)$$

Now denote by P the $d \times m$ matrix whose rows are $\phi_1, \phi_2, \dots \phi_d$. Then the approximant to x is $P^*P(x - \bar{x}) + \bar{x}$, and $y = P(x - \bar{x})$ are the new coordinates for x in the new, d -dimensional subspace.

In this derivation, we have assumed an *a priori* known, fixed, d , but in practice one constructs R , and then examines its eigenvalues, choosing d such that the error (3.3) is below whatever threshold one decides.

3.2 Proper Orthogonal Decomposition of quantum trajectories

In attempting to better understand, and approximate, open quantum systems, we might like to find lower-dimensional spaces in which the dynamics are confined, and for this we turn to Proper

Orthogonal Decomposition. Applying the POD process to quantum trajectories requires making a few, nontrivial choices. What space will we define as the larger space of which we seek the optimal subspace? What is the appropriate measure of distance between states, to evaluate the error and optimality of the subspace? A quantum trajectory, in order to be computationally tractable, requires measurement of every output, meaning that states remain pure at all times. We therefore start with wavefunctions, rather than density matrices. However, we know that the experimentalist will almost never have full observation of their system, meaning that she will be working with density matrices.

Density matrices have an amenable algebra, with a clear inner product

$$\langle \rho_1, \rho_2 \rangle = \text{Tr}(\rho_1 \rho_2) \quad (3.6)$$

for use in Eqn. (3.5). In addition, the form that a “direction in density matrix space” will take is clear: a trace-0, Hermitian matrix. For small perturbations of a density matrix by addition or subtract of such a matrix, we know the resulting matrix will almost always remain a valid density matrix — trace 1, positive, and Hermitian. We will see below that large perturbations can result in non-physical density matrices. In contrast, it is not clear how to define a “direction in wave function space.”

Of course, the POD algorithm requires its input to be vectors, in order to calculate the correlation matrix R . My process was to create the density matrix corresponding to each wavefunction, $\rho_i = |\psi_i\rangle\langle\psi_i|$, and then to “vectorize” each density matrix. Rather than simply creating a vector which was the end-to-end concatenation of each row or column of ρ_i , I took advantage of the Hermitian structure of the matrix to make a completely real vector by first concatenating the real parts of the rows of the upper-triangular portion of the density matrix, and then concatenating that with the imaginary parts of the upper-triangular portion (aside from the diagonal, which is entirely real). This significantly simplified the computer code necessary to implement the algorithm, and eliminated the chance of error due to confusion of which transpose operations should be complex-conjugate transpositions, and which should be simple transpositions. The * operations in the derivation above are complex conjugate transpositions, but derivations in the literature are not specific on this point, as they assume that everything is simply real. In practice I verified that if one simply unwraps the density matrix, leaving its vector form complex, and treats these as complex-conjugate transpositions, one achieves the same results as I do with the less-mistake-prone real-only method.

I chose a simple Itô-Euler integrator for the propagation of quantum trajectories, as well as for filters. The Quantum Optics Toolbox [10] provides two integrators, an Euler integrator and CVODE. CVODE uses a variable time step, which is incompatible with saving the full homodyne photocurrent record for reuse with the filter. Using the same integration algorithm, with the same time step, is essential for proper functioning of the filter. I used only homodyne detection of the

cavity field and atom because the heterodyne detection portion of the Quantum Optics Toolbox malfunctions: the mean state for a long trajectory is not equal to the steady state solution of the master equation. For the long trajectories (in both phase and absorptive bistability regimes) in this thesis, I chose a small time step, $\Delta t = 5 \times 10^{-6}$, but for shorter trajectories in the phase bistability regime, a time step of as large as 10^{-5} was possible without breeding instability. For absorptive bistability, $\Delta t = 10^{-4}$ was possible for shorter runs. For the test runs for the filters, $\Delta t = 2 \times 10^{-5}$ for absorptive bistability and $\Delta t = 2 \times 10^{-6}$ for phase bistability (which has much faster dynamics than absorptive bistability). The long trajectories took 3 to 4 days to run, with the POD processing taking several hours.

Rather than using the wavevector at every time step for POD, I sampled every N time steps, where N was generally between 100 and 1000 depending on Δt . The time to run the POD algorithm scales linearly with the number of wavevectors used. In order to stay within the memory limitations on the computer (and within the file size limitations of the *matio* C library), I saved the wavevectors in files of roughly 1000. This meant initializing the integrator for many separate, shorter runs, with the initial state of each corresponding to the final state of the previous. The Quantum Optics Toolbox cannot output both operator expectation values and the wavevector from a single run, so when I needed both (such as for the shorter runs used to test the filter), I was forced to run each time slice twice, with the same random seed. The Proper Orthogonal Decomposition algorithm was written in a combination of C and Matlab. I used the *matio* library written by Kevin McHale to read and write Matlab files from C, and took advantage of the BLAS linear algebra library for the matrix multiplication to calculate the correlation matrix R . BLAS can take advantage of the multiple processors on our server, resulting in significant speedup.

This definition of a vector form \vec{v} of ρ is also consistent with the 2-norm for which the optimality of POD is proven:

$$\begin{aligned} \|\rho_1 - \rho_2\|^2 &= \langle \rho_1 - \rho_2, \rho_1 - \rho_2 \rangle = \text{Tr}((\rho_1 - \rho_2)^2) \\ &= (\vec{v}_1 - \vec{v}_2) \cdot (\vec{v}_1 - \vec{v}_2) = \langle \vec{v}_1 - \vec{v}_2, \vec{v}_1 - \vec{v}_2 \rangle = \|\vec{v}_1 - \vec{v}_2\|^2. \end{aligned} \quad (3.7)$$

As we will see below, however, this distance definition is not necessarily consistent with the accuracy of POD-based filters for particular observables we care most about (such as the expectation values of field quadratures).

The output of the POD algorithm is a set of basis vectors for a density matrix space (a set of trace-0 Hermitian matrices). These are generated from the eigenvectors of the correlation matrix R and each state is associated with an eigenvalue which indicates its relative contribution. We rank the basis vectors according to this eigenvalue, and (as a first cut) discard any state whose eigenvalue falls below 0.1% of the sum of all eigenvalues (the trace of R). For our systems, this commonly leaves

between 10 and 50 eigenstates. Further cuts are always done in the order of increasing eigenvalue, but may be based on the performance of the projected filter or a more subjective sense that the states below a certain eigenvalue look more like noise than states which would significantly contribute to a more accurate approximation.

Proper Orthogonal Decomposition (and various modifications thereof) can produce subspaces which look like they include the relevant parts of the system dynamics. However, we must evaluate them quantitatively. To do this, we project the system dynamics onto several of these subspaces, using the method derived in Section (2.2), and compare the dynamics indicated by the filter with the trajectory dynamics. In this chapter, I use only the normalized dynamical system; the un-normalized system performs slightly less well, and feels slightly less rigorous, due to the necessary renormalization at each time step. The un-normalized projected system also requires re-orthogonalizing the set of states to be used, because the mean state is not orthogonal to the other POD basis states. This in turn means that we lose some of the clarity from the ordering of the states by eigenvalue.

3.3 Phase bistability

The density matrix (and associated Q function) for a phase bistable cavity QED system has a simple structure — roughly, it is the superposition of two compact states with equal amplitude but opposite phase. Q function dynamics consists largely of a compact, roughly Gaussian, peak located at one stable point or the other, or in transition along the path between them. This superficially linear form gives us hope that the linear approximations made in the POD algorithm will not significantly impair performance, and a relatively low-dimensional, high fidelity approximation will be possible. In fact, this hope is largely fulfilled.

Our canonical phase bistable system, used for all of the examples in this thesis, has the following parameter values: $\Theta = 0$, $\Delta = 0$, $\kappa = 4$, $\gamma = 2$, $g_0 = 12$ and $\mathcal{E} = 23.57$. This corresponds to the set of parameters also used in [11]. (Van Handel and Mabuchi use superficially different values, but the ratios determine behavior; in effect I have simply scaled time differently). Expectation values from an excerpt of an example trajectory are shown in Figures 3.3 and 3.4. The trajectory I used to generate the basis states for the subspace onto which we hope to project dynamics was run for 5000 units of time, with a time-step of $\Delta t = 5 \times 10^{-6}$.

The Q functions of the mean (steady state) and the first 11 eigenstates (potential basis vectors) are shown in Figure 3.1. While Q functions for valid quantum states are by definition always positive, the potential basis vectors have Q functions which can swing wildly negative. Recall that these are not quantum states, but rather directions; the Q function is simply an aid to the eye and understanding. A projected state consists of some linear combination of these basis vectors, added to the mean state; the Q functions roughly “add” as well. For example, a state which is composed

by adding the leading eigenstate to the mean would add probability density to one of the two states, while removing it from the other; in effect this first basis direction alone can capture the switching behavior (as we will see in the filter performance evaluation below). Notice that this switching behavior will generally result in rather large components in the direction of this leading eigenstate, threatening the “small perturbation” limit in the linearizing POD process. I get away with this stretching of the POD approximations, I believe, only because the phase bistable system is so simple and symmetric; if the two states were asymmetric I would be forced to undertake the additional complicating procedures I discuss below in the context of (asymmetric) absorptive bistability.

The subsequent states capture behavior within the switching region, or dynamics within one state. The transition basis states take a form similar to sinusoidal peaks and valleys between the two stable states, and we may think of them as the basis states for a Fourier decomposition of a single peak traveling between the two states (behavior which the leading eigenstate cannot capture). Eigenstates beyond the 12th become hard to understand, and their eigenvalues are all less than 0.5% of the sum of all eigenvalues. If we keep the first 12 states, we have kept eigenvalues which sum to 97.3% of the total of all eigenvalues. See Figure 3.2 for a plot of the 60 largest eigenvalues in descending order.

3.3.1 Projection filters

In order to evaluate the performance of the POD-produced approximations, I ran a long, high-resolution (small time-step), quantum trajectory, separate from the trajectory on which the POD algorithm was run. The homodyne photocurrent from the phase quadrature measurement, as well as the expectation values of many system operators, were recorded for comparison with the filter. In addition, I kept 5000 sampled wavefunctions in order to evaluate the fidelity of the approximate states using a variety of measures of quantum fidelity.

Figure 3.3 shows the expectation values for the phase quadrature of the cavity field, $\langle y \rangle$, from a short excerpt of this trajectory, with the quantum trajectory in blue. Comparable traces of $\langle y \rangle$ from filters produced by projecting onto a subspace with 1, 4, 8, and 12 basis states are shown in green, red, cyan, and magenta. Examining these in order, we first see that even a single-dimensional subspace captures the switching behavior of the system (and it does it while stretching the definition of a “small” difference from the mean implicit in our linearized system). However, the 1-D subspace does not capture the transition dynamics between the two states — on the second transition in Figure 3.3, it is closer to a square wave than the decaying exponential of the full trajectory. Adding three basis states, which include some information about states between the two stable ones (see Figure 3.1), allow this transition to exhibit some “decay,” while with 8 or 12 states the approximate system can follow the exponential more closely.

Looking now at the first transition, the instabilities of the filter system become apparent. A

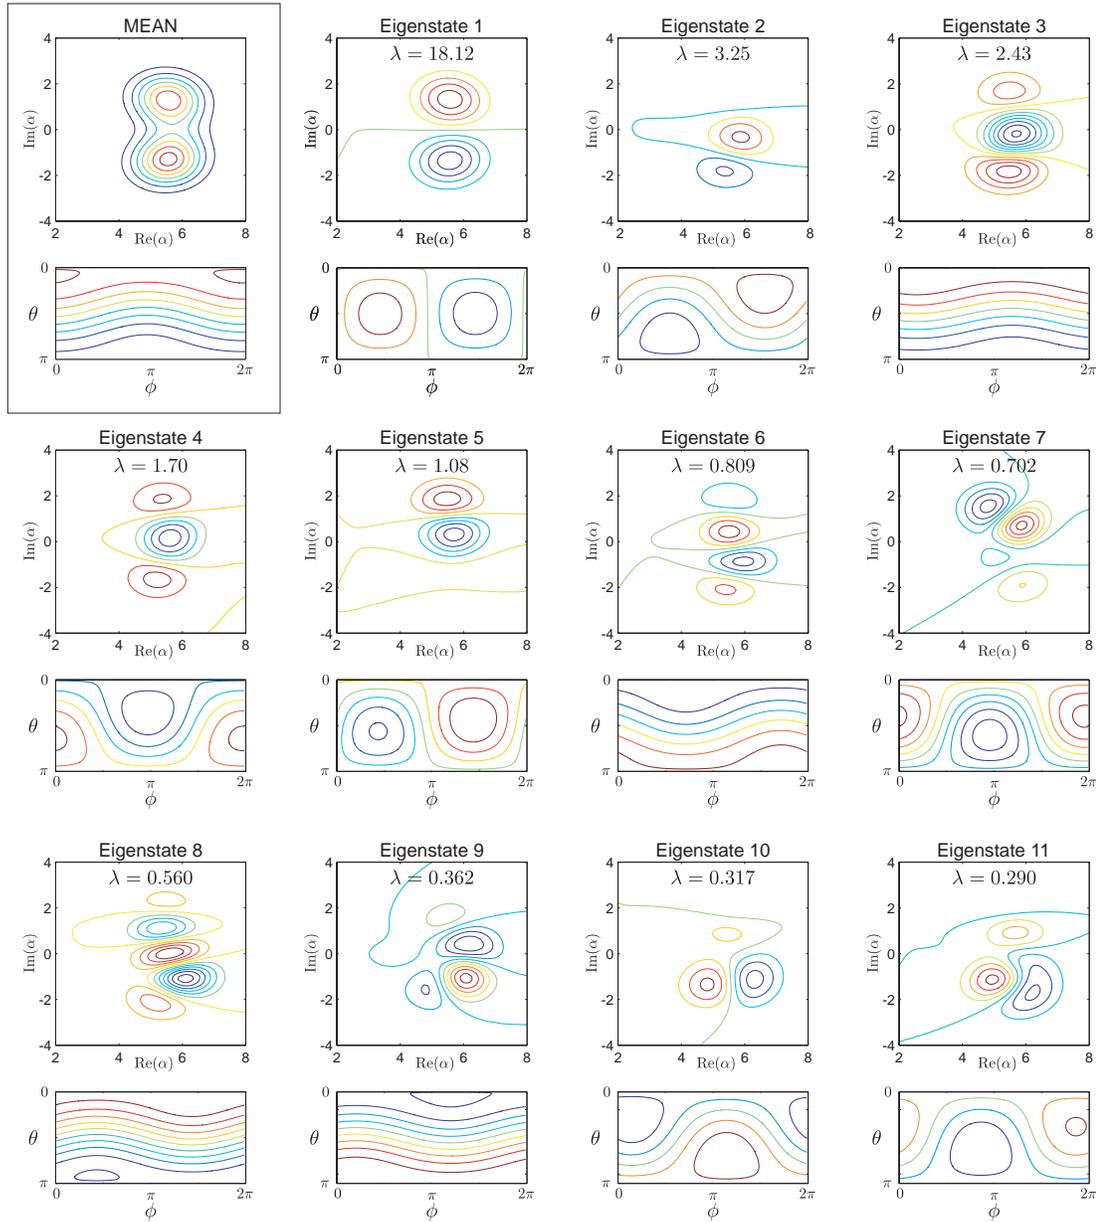


Figure 3.1: Contour plots of Q functions for the mean and the first 11 orthonormal eigenstates from a long phase bistable trajectory. The upper plot for each state is the Q function of the field state after tracing over the atom; the lower plot is the Q function of the atom after tracing over the field (produced using Generalized Coherent States [33]). The eigenvalue corresponding to each eigenstate is included, showing the relative importance of each basis state/direction. The first eigenstate is responsible for most of the switching behavior, whereas the others are largely responsible for motion within each eigenstate (basis states 2, 5, 7, 10 and 11) or the transitions between states (basis states 3, 4, 8, and 9). The atomic Q functions are parametrized on a sphere spanned by $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$. Eigenstates 3 and 4 have very similar cavity field effects, and roughly orthogonal atomic states.

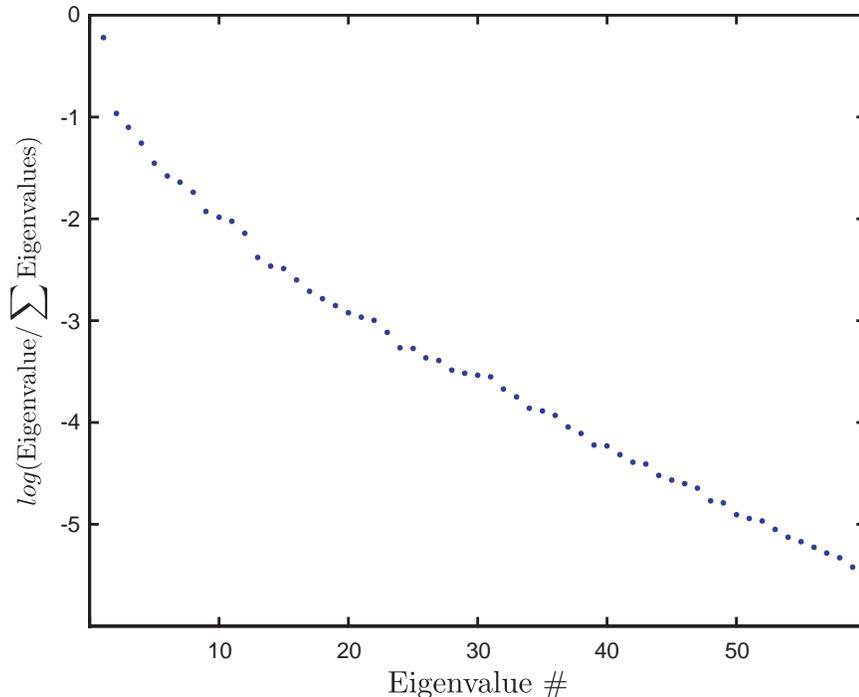


Figure 3.2: Eigenvalues of the correlation matrix R for the canonical phase bistable case, shown as the logarithm of the fraction of total eigenvalues, for the 60 eigenstates with largest eigenvalues.

segment of (noisy) photocurrent here apparently contains relatively little accurate information about the state, the state is only very poorly approximated by the limited subspaces in these 12-or-fewer dimensional approximations, or the photocurrent excites the use of basis states which are not helpful for stability. Similar patches of instability occur with some regularity, and are worst with 12 basis states. In these areas, the filtered state is no longer a valid quantum state (with a positive semi-definite density matrix). Not surprisingly, the fidelity measured between the wavefunction from the quantum trajectory and the projected filter is quite poor in these areas.

I used two fidelity measures to evaluate the performance of this filter: Probability of Error and the Bhattacharyya Coefficient. See Fuchs and van de Graaf [34] for a derivation of these two coefficients, as well as the Kolmogorov Distance (which is intimately related the Probability of Error and doesn't contain any more information). The Probability of Error is defined as

$$PE(\rho_0, \rho_1) = \frac{1}{2} - \frac{1}{4} \sum_{j=1}^N |\lambda_j|, \quad (3.8)$$

where λ_j are the eigenvalues of $(\rho_0 - \rho_1)$. This value should vary between 0 (two states are perfectly distinguishable) and 1/2 (perfectly indistinguishable) as long as ρ_0 and ρ_1 are valid density matrices.

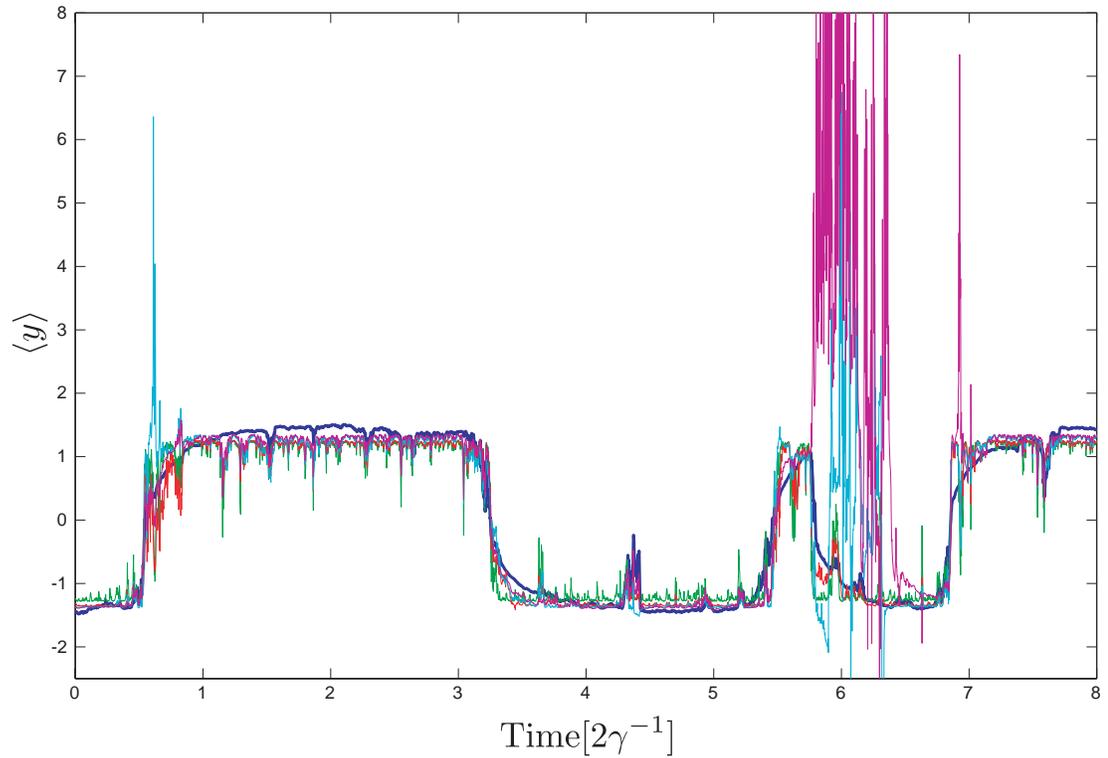


Figure 3.3: A short excerpt from a single phase-bistable quantum trajectory, showing the expectation value of the phase quadrature of the cavity field, $\langle y \rangle$, and the trajectories of three filters using the same homodyne detection photocurrent. The trajectory is in blue, and the projected filters using 1, 4, 8, and 12 basis states are shown in green, red, cyan, and magenta, respectively. The filters with 8 and 12 basis states are somewhat less stable (witness the first and fourth transitions here), but give a closer approximation during the “quiescent” times between transitions, and fit the exponential decay shape of the second transition better than the 1 or 4 basis-state approximations. Projecting onto even a one-dimensional subspace, however, successfully captures the switching behavior of the system.

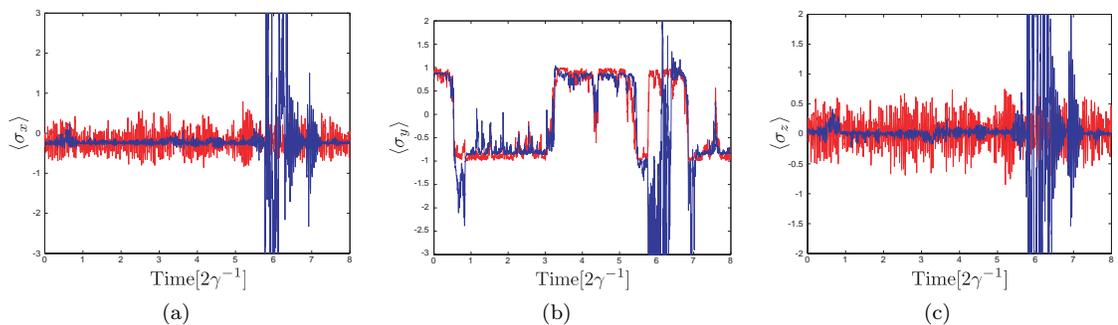


Figure 3.4: Trajectories of the three simple atomic Hermitian observables for the same time span as Figure 3.3. (a) $\langle \sigma_x \rangle$, (b) $\langle \sigma_y \rangle$, and (c) $\langle \sigma_z \rangle$. For each observable, the red trace is the quantum trajectory; the blue trace is the filter with 12 basis states (the same filter as shown in magenta in Figure 3.3).

The Bhattacharyya Coefficient is defined as

$$B(\rho_0, \rho_1) = \text{Tr} \sqrt{\sqrt{\rho_0} \rho_1 \sqrt{\rho_0}}. \quad (3.9)$$

Its value varies between 0 (orthogonal states) and 1 (identical states).

The calculated values of these two fidelity measures over the course of an extended filtering simulation offer some surprises, which indicate both the failures of the filter, and differences between these two measures. The average Probability of Error decreases as we include more states in our basis, from 0.2065 with one basis state, to 0.2059 with 4, 0.1057 with 8 and -0.0708 with 12. What is happening is that the periods of instability are dominating the average. In these areas, the approximate state is not a valid quantum state, and the value of the “probability” can fall below zero. However, if we examine just the quiescent periods between state transitions and unstable patches, the Probability of Error is almost identical for the four filters.

The Bhattacharyya Coefficient (BC) behaves somewhat differently. In particular, when the filter state becomes unphysical, the BC becomes imaginary, rather than negative (I use only the real part of the BC, as the mean of the imaginary parts is very close to zero). This allows the average value over a filtering run to stay almost constant with the size of the basis state, because the 12-dimensional basis does provide for a slightly higher BC than the 8 in the “normal” periods, which in turn has a higher BC than the 4- and 1-state filters. The unstable zones have a BC of zero, so the increased instability from added basis states results in a falling average. The 1 state BC is 0.4364, 4 state BC is 0.4485, 8 state is 0.4463, and 12 state is 0.4355.

Proper Orthogonal Decomposition appears to work quite well to define a subspace in which the phase bistable dynamics take place. The basis states (in Figure 3.1) capture both stable states, and allow for approximations of both the dynamics within each state and the transitions between. When stable, the projected filter can do a good job of matching the behavior of $\langle y \rangle$ and $\langle \sigma_y \rangle$; other observables fluctuate quickly, while the filter approximates their mean. The occasional instability of the filter, however, raises doubts as to whether POD is a robust solution.

3.4 Absorptive bistability

Absorptive bistability occurs as part of a much more complicated system than phase bistability. In particular, the upper “state” is extended for our set of parameters (very different from the close-to-minimum-uncertainty states in phase bistability). Complicating the performance of the Proper Orthogonal Decomposition algorithm, the system can be significantly asymmetric between the upper and lower states, depending on the driving field strength (phase bistability is structurally exactly symmetric). I have chosen a driving field for the examples in this thesis at which the fraction of time

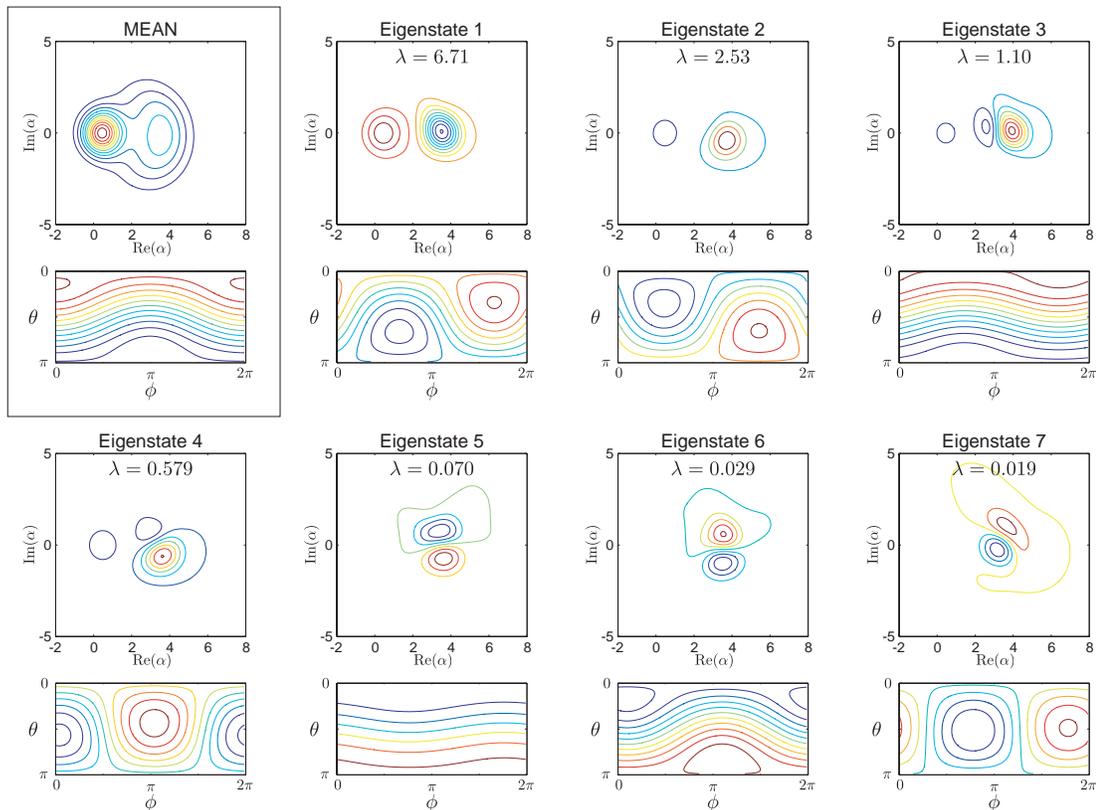


Figure 3.5: The mean and leading 7 eigenstates from the Proper Orthogonal Decomposition algorithm applied to a long absorptive bistable trajectory. The first eigenstate allows the filter to remove almost all traces of the lower state, while the remaining eigenstates finish that cancellation and allow for dynamics within the upper state. No combination of these states, however, allows for the cancellation of the full upper state in order to accurately approximate the lower stable state. The model parameters are $\Theta = 0$, $\Delta = 0$, $\kappa = 0.1$, $\gamma = 2$, $g_0 = \sqrt{2}$, and $\mathcal{E} = 0.56$, and the simulation was run for $t_{final} = 5000$ (with time scaled by $\gamma/2$).

spent in each state is roughly equal; however, this may not be the experimentalist’s preferred (or available) parameter value, resulting in lower POD performance unless the algorithm is extended, as I will describe later in this subsection.

The parameters I have chosen to demonstrate the application of POD to an absorptively bistable system are: $\Theta = 0$, $\Delta = 0$, $\kappa = 0.1$, $\gamma = 2$, $g_0 = \sqrt{2}$, and $\mathcal{E} = 0.56$. (This is clearly a “good cavity” relative to our phase bistable system; we achieve similar cavity field amplitudes for almost 2 orders of magnitude less driving field intensity.) Notice that this driving field strength is above the upper semi-classical (Maxwell-Bloch) bifurcation point ($\mathcal{E} \approx 0.553$); that is, the semi-classical system has only one stable point at this driving field value. These are similar to the parameters chosen by [22], except I have chosen a slightly lower driving field in order to get closer to an even balance between the lower and higher amplitude states.

Simply applying the POD algorithm to a long ($t_{final} = 5000$) trajectory (with homodyne measurement of field amplitude quadrature) results in the set of states shown in 3.5. Note that the

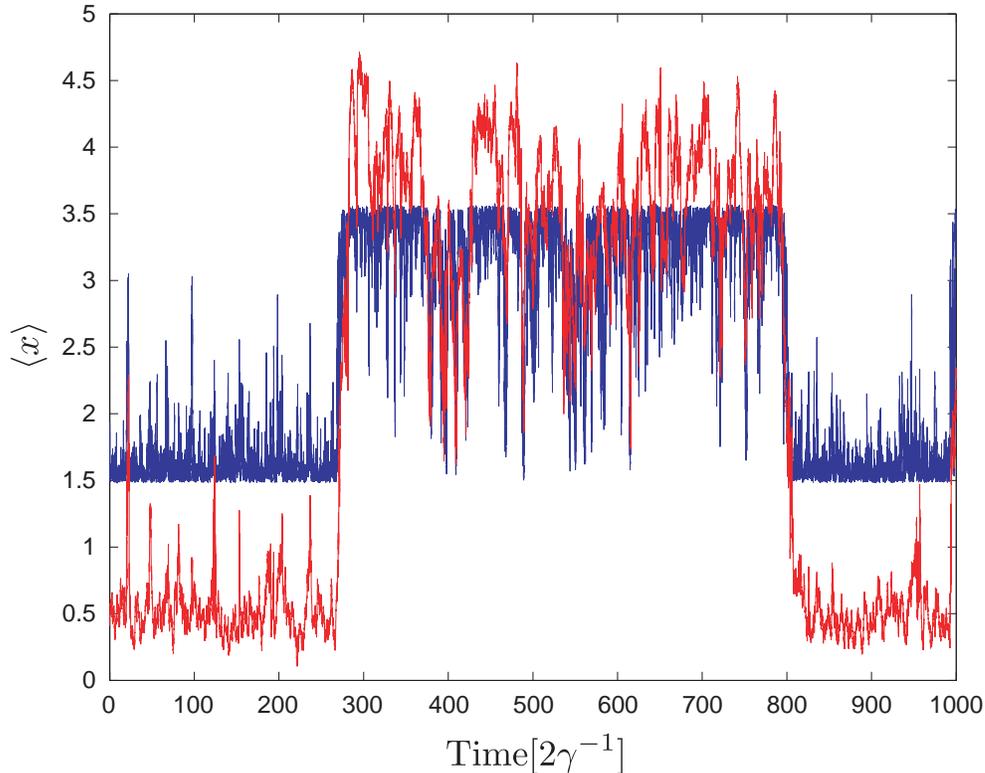


Figure 3.6: Comparison in the absorptive bistability regime of simple POD filter with exact quantum trajectory. The blue trace is the expected value of the amplitude quadrature for a POD filter with a 7-dimensional linear subspace; the red trace is from the quantum trajectory. The eigenvalues corresponding to these leading 7 eigenstates sum to 99.85% of the sum of all eigenvalues; the largest eigenvalue left out is less than 0.08% of the total. Amplitude quadrature homodyne measurement; params.

states beyond the first generally relate to directions in the upper state, and none of these first 8 basis vectors allow for dynamics or fluctuations within the lower state. We know from observing the trajectory that what dynamics there are within the lower state are smaller in amplitude than those within the upper state. The first eigenstate is generally responsible for the switching behavior: subtract this from the mean to move towards the upper state alone, and add it to move toward the lower state alone. Note, however, that the shape of the upper portion of this basis state does not match the shape of the upper portion of the mean, so that adding this state alone will not put us exactly in the lower state. The shapes match much more clearly for the lower state, so if we subtract the correct multiple of this basis state from the mean we can almost exactly cancel the lower state. The second and higher basis states allow for some fluctuations within the upper state.

This asymmetry between the lower and upper states results in significantly impaired performance as a filter. Using 7 eigenstates to make the basis of the subspace into which we project the dynamics, we plot $\langle x \rangle$ from the projected dynamics in Figure 3.6. The eigenvalues corresponding to these leading 7 eigenstates sum to 99.85% of the sum of all eigenvalues; the largest eigenvalue eliminated

is less than 0.08% of the total. The filter tries to match the field amplitude of the lower state, but simply cannot, because the subspace does not include a good approximation to the lower state. The mismatch in shapes results in a minimum achievable $\langle x \rangle$ which is above that of the lower state. The filter does relatively much better with the upper state, although it cannot capture all of the volatility. It is able to track some of the fluctuations down from the upper state, but cannot go higher than a fixed value. It is also unable to track the volatility in $\langle y \rangle$ within the upper state (not shown).

Needless to say, the fidelity of this filter is quite poor. The mean Probability of Error is negative (-0.008), thanks to long periods (in the lower state) in which the filter state is not a valid density matrix (negative eigenvalues, and the Q function not positive everywhere). The Probability of Error measure is positive during the times when the system is in the upper state, averaging around 0.02, with a maximum of 0.155. The mean Bhattacharyya Coefficient is 0.185, with an average of about 0.13 in the lower state, and 0.24 (but wildly varying between 0.012 and 0.6) in the upper state. The resilience of the Bhattacharyya Coefficient to unphysical filter states makes it a somewhat preferred measure.

Some trajectories randomly spend a larger fraction of their time in the lower state than the one I used here, and these trajectories can exhibit the opposite problem as a filter: they capture the lower state and its dynamics well, but badly miss the upper state. In order to address this problem, I tried two different strategies: weighting the sum used to calculate distances in the density matrix space, and dividing the data into three “zones”: lower, transition, and upper, and executing POD separately in each zone.

3.4.1 Weighted POD

The observables we tend to care most about in evaluating the accuracy of our model reduction are the 5 observables whose dynamics can be approximated by the Maxwell-Bloch equations: $\langle x \rangle$, $\langle y \rangle$, $\langle \sigma_x \rangle$, $\langle \sigma_y \rangle$, and $\langle \sigma_z \rangle$. In matrix form (to act on density matrices), the nonzero terms for these observables lie on the diagonal or immediate off-diagonal. (For a tensor product system like an two-level-atom and cavity system, let us define the “diagonal” to be the diagonal of each quadrant.) The values of these observables (calculated as $\text{Tr}(\mathcal{O}\rho)$ for each operator \mathcal{O}), then, depend only on the diagonal and immediate off-diagonal terms in the density matrix. Proper Orthogonal Decomposition, on the other hand, treats every entry in the density matrix identically. POD produces the optimal subspace provided that we share its definition of distance, which treats the real and imaginary parts of each entry in a completely even-handed fashion. The failure of POD to produce a single subspace which includes both the upper and lower states can be partly attributed to this evenhandedness: dynamics within the upper state win out over including an accurate approximation for the lower state.

To address this, I tested various re-definitions of distance by weighting the terms of the density

matrix. To do this, I defined the distance from the diagonal n_d to be the absolute value of the difference between i and j labels for each entry in the matrix (modulo half the matrix size to allow for the tensor product structure). I then divided each entry by n_d^X where $X = 0.5, 1, \text{ or } 2$. This process treats the diagonal differently from the immediate off-diagonal, so I also used a weighting scheme which measured distance from the immediate off-diagonal, so that the weighting would follow the pattern (B) $1, 1, 1/2, 1/3, 1/4, \dots$ instead of (A) $1, 1/2, 1/3, 1/4, \dots$ as we proceed away from the diagonal. None of these six weighting schemes has any appreciable effect on the balance between the upper and lower states, relative to basic POD ($X = 0$). All produce basis states similar to those shown in Figure 3.5, with only dynamics of the upper state represented, and fail as filters because they cannot come close to approximating the lower state.

There is another inequity between the distribution of density matrix entries and the distribution of expectation values: the density of Fock states scales like the square of the amplitude of a coherent state, so that a state further from the vacuum has large entries in many more locations in the density matrix than one close to the vacuum. Imagine two pairs of states, one pair close to the vacuum, and one with $\langle x \rangle \approx 4$. The distance within each pair, when measured in the way we defined for Proper Orthogonal Decomposition, is the same. However, the pair close to the vacuum will have substantially different expectation values, while the higher amplitude pair might be almost indistinguishable by that measure. What skews POD is the reverse of this: small fluctuations in expectation values at high amplitude look like large distances in the POD metric, while small fluctuations near the vacuum are small in this metric. This is compounded by the behavior in the absorptive bistability regime, which has larger fluctuations in the upper state, and the POD process outputs states which focus on the upper state and do a very poor job in the lower.

To counteract this, I tried a weighting scheme which weights the vacuum state highest, and decays away from there. We know *a priori* that the scaling should go like the square root of the Fock state in order for the distance to be independent of the coherent amplitude. Applied to the POD process, this weighting also has no appreciable effect on the fundamental problem: the lower state is not accurately approximated by combinations of the dominant eigenstates. This makes sense because at the same time that this weighting reduces the effective distance between density matrices within the upper state, it also undercounts the remnants of the upper state which are left behind when using the leading eigenstate (see Figure 3.5) to attempt to approximate the lower state. For weighted POD to improve on the unweighted process, we would need a weighting scheme which somehow simultaneously counts and discounts the higher-field-amplitude portions of the density matrix. One possibility, unexplored in this thesis, would be to move the origin for the POD analysis to the space between the lower and upper “states,” so that they both have similar amplitude and take up similar “space” within the density matrix. They will still be asymmetric, however, in both shape and the time the system spends in each (especially if the driving field is not selected for balance between

states). Fundamentally, Proper Orthogonal Decomposition with a single mean, and perturbations around that mean, has difficulty handling dynamics around two separated, asymmetric states.

3.4.2 Zoned POD

Proper Orthogonal Decomposition implicitly makes “small perturbation” assumptions, which we violate when we consider a system with multiple stable points, as the leading eigenstates end up with quite large coefficients to push onto into one stable regime or another. In addition, the algorithm’s emphasis on the dominant, most common states means that it tells us almost nothing about the transitions between states when there are significant dynamics within states. The phase bistable system considered in Section 3.3 is symmetric enough, and switches often enough, that the switching dynamics and the transition periods are on a more level footing with fluctuations and dynamics within each state. The absorptive bistability regime is not symmetric, and transitions are much more rare. We can improve compliance with assumptions, capture some transition behavior, and re-balance asymmetry by dividing the data into *zones* surrounding each stable point, and in between them.

For the driving field strength, cavity damping rate, and other parameters which define our standard absorptive bistability regime, the lower state has a mean cavity field amplitude of about 0.5, and the upper state (while broader) has a mean of about 3.25. The fluctuations around these two states rarely cross into the range of 1.5 to 2.5 unless the system is actually in transition between the two stable regions. Therefore, I cut the data into three sets: “lower” for any point in the trajectory which has $\langle x \rangle < 1.5$, “upper” for points with $\langle x \rangle > 2.5$, and “transition” for points in between. (I include a short segment on either side of the exact transition zone to associate behavior leading to or out of transitions with them, rather than with the lower or upper zones.) Note that the constant measurement and decay of the system helps keep the field states compact. They are almost always roughly Gaussian (with a Q function width similar to the width of a coherent state), so we can use $\langle x \rangle$ as a proxy for the full distribution.

Once we have divided the data into three subsets, we have two options for defining the mean state for the POD algorithm: we could continue to use the overall mean, or find the mean of each zone. In the former case, the leading eigenstate will be the direction from the overall mean to some part of that zone, and subsequent basis states provide a direction within the zone. The latter case creates the transition direction directly, without use of POD, simply by subtracting the zone’s mean from the overall mean. All the POD basis states are then directions away from that local mean. I have tried both definitions, but selected the latter because it is most consistent with the small-coefficient assumptions built in to POD and because the means will be independent of small asymmetries which might otherwise throw the leading eigenstate “off balance.” This is particularly a concern for the transition zone, where the relatively small number of transitions means that if more transitions take

place with $\langle y \rangle > 0$ than $\langle y \rangle < 0$, the “one mean” algorithm would likely select the leading eigenstate to reflect just the $\langle y \rangle > 0$ transitions, biasing all the subsequent eigenstates as well.

The output of three separate POD processes (one for each zone) is three sets of Hermitian matrices, from which we will build our subspace. (Implementation of the algorithm tends to produce matrices which are almost but not exactly Hermitian, due to numerical errors; I added a simple step which tweaks them to be exactly Hermitian.) There are four matrices with trace 1 (valid density matrices): the overall mean and the mean of each zone. I turn the zone means into trace-0 (direction) matrices by subtracting the overall mean. The remaining matrices are already directions in density matrix space, but they are not all orthonormal. The directions within each zone are orthogonal (although not orthogonal to the newly-created “mean direction” for that zone), but they are not orthogonal between zones. We therefore run the set of all trace-0 direction matrices (including the means) through a simple inner-product based orthonormalization process. Because they are all Hermitian, their inner products are all real, and their linear combinations with real coefficients produce new, orthonormal, direction matrices which remain Hermitian. The linear combinations which result from filtering will therefore be valid density matrices as long as they remain positive definite.

Figures 3.7 and 3.8 show the leading states from two different trajectories in the absorptive bistability regime. Figure 3.7 results from a trajectory in which the cavity field was measured in the amplitude quadrature; Figure 3.8 was measured in the phase quadrature. (Note that both show their set of states before they have been orthonormalized.) Because the steady states and behavior of a quantum system do not depend on the measurement, it’s not surprising that these two sets of states are quite similar. They do provide an opportunity to examine some of the kinds of differences which result from different trajectories. For example, the lower zone eigenstates are arranged at different angles, indicating that the noise in the lower state scattered more in one axis than another during the trajectory.

In both cases, the transition zone is biased towards $\langle y \rangle > 0$ behavior in its leading eigenstates. Each of these trajectories include only a dozen or so transitions, meaning that the statistics are quite poor. A much longer trajectory (which would require a more patient researcher and/or a more powerful computer) ought to have more balanced eigenstates in the transition zone. These higher-quality eigenstates might allow us to see whether transitions favor particular paths in phase space. In particular, the semi-classical Maxwell-Bloch equations have an unstable point between the two stable equilibria (see Figure 1.2), and these transition states might reveal whether the quantum system avoids the $\langle y \rangle = 0$, $\langle x \rangle \approx 2$ area, or not. I cannot draw conclusions on this point from the limited trajectories I have been able to analyze.

To test the performance of these POD subspaces, I ran them as a filter on the same homodyne photocurrent as in Figure 3.6 and the surrounding discussion. A disadvantage of the zoned approach

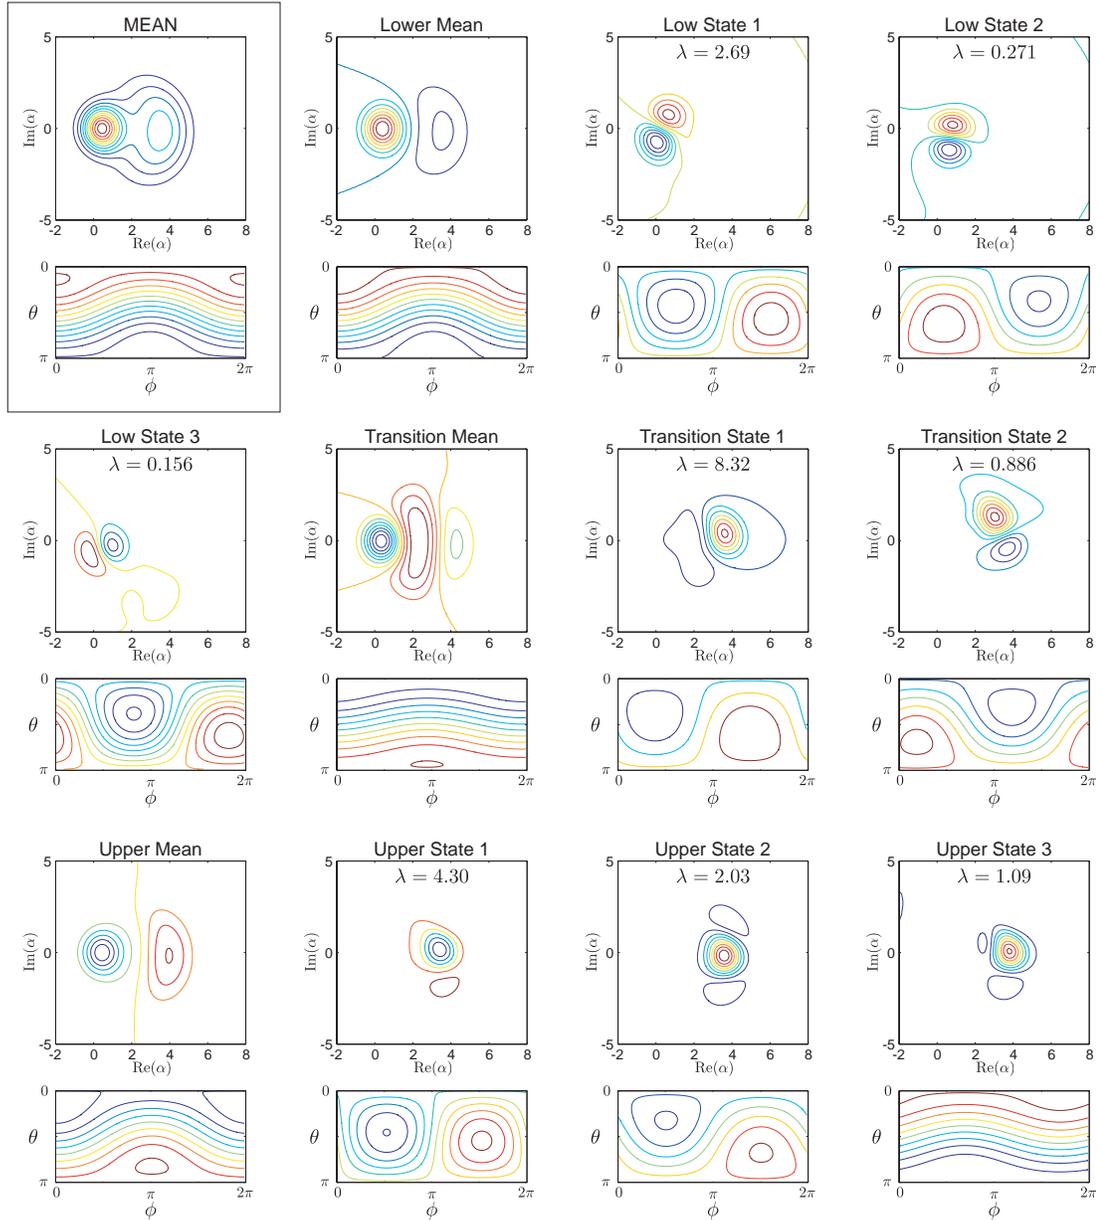


Figure 3.7: Contour plots of Q functions for states and basis vectors from “zoned” Proper Orthogonal Decomposition of absorptive bistability regime, with homodyne measurement of cavity field *amplitude* quadrature. The upper plot for each state is the Q function of the field state after tracing over the atom; the lower plot is the Q function of the atom after tracing over the field (produced using Generalized Coherent States [33]). Four states are generated separately from POD: the overall means state and the mean states in each of three zones. The leading 2 or 3 eigenstates from each zone are shown, along with the eigenvalue corresponding to each, showing the relative importance of each basis state/direction. The atomic Q functions are parametrized on a sphere spanned by $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$.

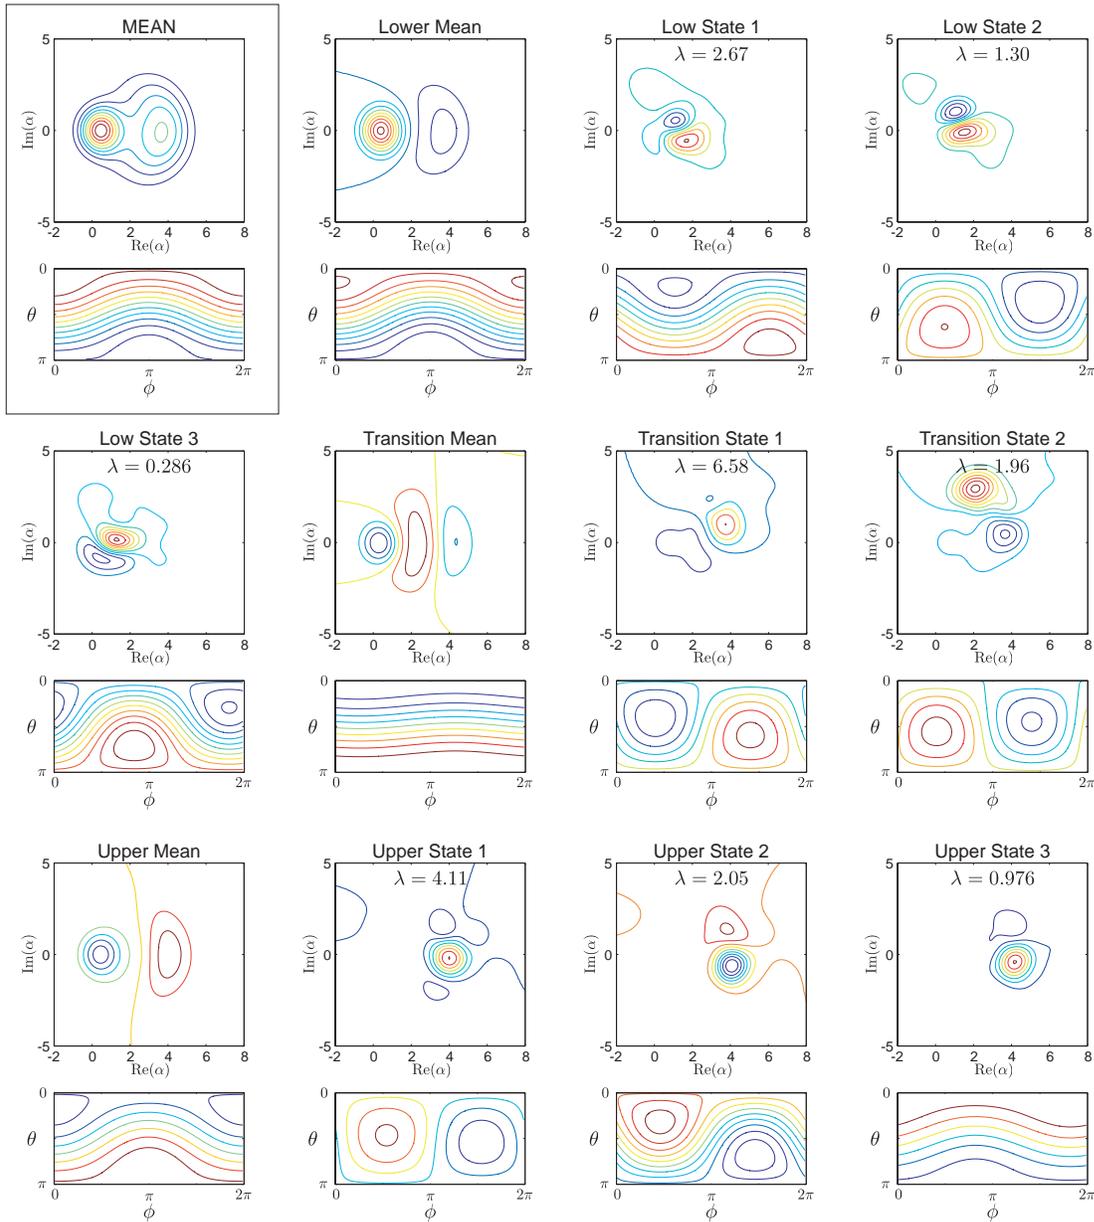


Figure 3.8: Contour plots of Q functions for states and basis vectors from “zoned” Proper Orthogonal Decomposition of absorptive bistability regime, with homodyne measurement of cavity field *phase* quadrature. The upper plot for each state is the Q function of the field state after tracing over the atom; the lower plot is the Q function of the atom after tracing over the field (produced using Generalized Coherent States [33]). Four states are generated separately from POD: the overall means state and the mean states in each of three zones. The leading 2 or 3 eigenstates from each zone are shown, along with the eigenvalue corresponding to each, showing the relative importance of each basis state/direction. The atomic Q functions are parametrized on a sphere spanned by $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$.

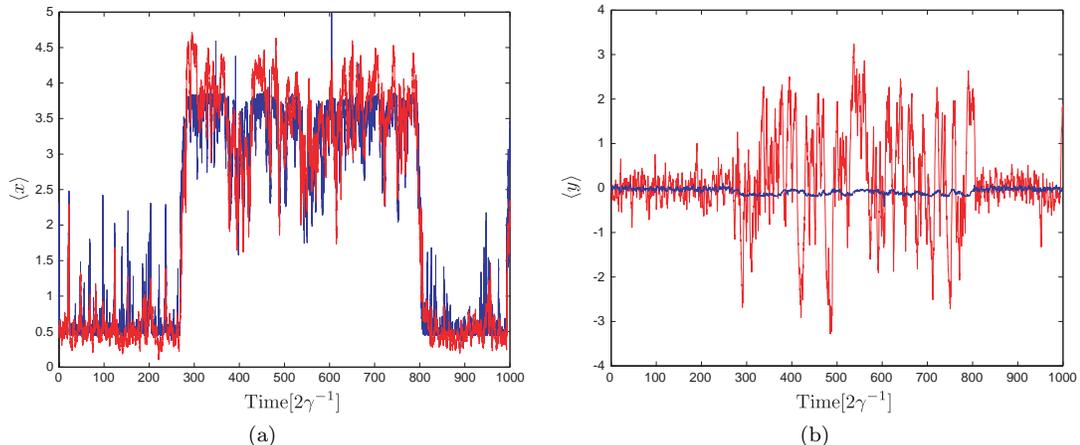


Figure 3.9: Comparison in the absorptive bistability regime of a “zoned” POD filter with exact quantum trajectory. The blue trace is the expected value of the amplitude quadrature for a POD filter with a 12-dimensional linear subspace; the red trace is from the quantum trajectory. The subspace is constructed from the mean states within each zone and the leading 3, 2, and 4 eigenstates from the lower, transition, and upper zones, respectively. Filter constructed and run on a system with a homodyne measurement on the amplitude quadrature.

is that we need to include more basis states, because we need the leading few from each zone. As an example, I chose to run the filter using a 12-dimensional subspace: the “lower mean” and the leading 3 eigenstates from the lower zone, “transition mean” and leading 2 eigenstates from the transition zone, and “upper mean” and 4 states from the upper zone. This means I am including 98.9% of the eigenvalue sum in the lower zone, 91.2% in the transition, and 98.3% in the upper zone, with 12 dimensions to consider (instead of 7 as in the basic POD filter above). As expected, this filter significantly outperforms the basic POD filter, helped immensely by its ability to successfully approximate both the lower and upper states. Figure 3.9 demonstrates the ability of the amplitude quadrature zoned filter to successfully capture the lower and upper states, as well as some of the dynamics within each. The large fluctuations within the upper state still prove to be too much for the “small amplitude” approximation built into the POD formulation.

The Probability of Error fidelity measure is somewhat better for the “zoned” POD than for the basic algorithm: at least the mean is positive (0.0134). The mean in the lower state is approximately 0.005, and it is somewhat better (0.027) in the upper state away from the few points where the filter becomes unphysical. The maximum Probability of Error for this data set was 0.09. The Bhattacharyya Coefficient similarly shows slight, but notable, improvement using the “zoned” algorithm relative to the basic POD process. The mean value of BC for this data set is 0.199 (compared to 0.185). The average in the lower state is about 0.17, and 0.22 in the upper state. Note that this reflects the markedly better performance for the zoned algorithm in the lower state and slightly worse performance in the upper state (perhaps due to the inclusion of only 4 eigenstates related to dynamics within that zone, compared with 6 in the simple POD algorithm).

If we discard the results of POD entirely, and use only the mean states for each zone, we end up with a reasonable filter, at least for amplitude quadrature behavior (it has no flexibility to capture the phase quadrature or atomic behavior). This 3-dimensional filter underperforms the 12-dimensional one, but might suffice for some applications. The mean Probability of Error is 0.0086 (with a mean of 0.0058 in the lower state, 0.0192 in the upper state, and a maximum of almost 0.08). The Bhattacharyya Coefficient is almost unchanged in this case for the lower state (where internal dynamics play a smaller role, so discarding them has a smaller effect), with a mean of 0.17. It performs less well in the upper state, as expected, with a mean of 0.215. The overall mean BC for this non-POD filter is 0.195.

The filter for homodyne measurement of the phase quadrature, using the same distribution of eigenstates as the amplitude quadrature, utterly fails as a reasonable filter. It behaves somewhat reasonably in the lower state, but fails in the upper state. To see why, look at the variation of the phase quadrature, $\langle y \rangle$, in Figure 3.9b, and compare it with the y -extent of the eigenstates in Figure 3.8. A filter based on the homodyne photocurrent from a phase quadrature measurement will attempt to match the observed value of the phase, but to do so given the eigenstates even from zoned POD requires unphysical states from combinations with large coefficients. In turn, these unphysical states do not do a good job of reproducing the other observables (like $\langle x \rangle$), and the fidelity of the filter to the actual dynamics is negligible (and the filter tends to be very unstable). This is not only a product of this “zoned” POD process: the same thing happens with basic, un-zoned POD, and with zoned basis states calculated using only the single overall mean (rather than the means for each zone).

3.5 Discussion

As we have seen in this chapter, Proper Orthogonal Decomposition is capable, when used with care and extended when necessary, of producing low-dimensional linear models which capture some of the dominant dynamics of quantum systems. Basic POD, applied to the phase bistable case, produced filters that are usually highly accurate at capturing the broad dynamics of the system, and might serve as reasonable filters for a control application. When extended into “zones,” POD also works quite well for absorptive bistability, as long as the basis states used include close to the full range for the variable being measured (the amplitude quadrature for our example).

POD is not a universal tool, however, and its limitations in both the phase and absorptive bistability regimes illustrate broader issues that would arise should the tool be extended to other systems. The first of these is the algorithm’s tendency to deliver spaces which include only the type of dynamics in which the system spends the largest fraction of the time. In the phase bistable case, the particular trajectory used as an example in this thesis spends slightly more time in one

state than the other. As a result, the basis set returned by POD favors some dynamics within this slightly-favored state over dimensions that inform us about state transitions; this may not match the order in which the experimentalist would emphasize these two behaviors.

This sensitivity to asymmetry is also at play in the absorptive bistability example, and is exacerbated by shape asymmetry between the upper and lower states. The definition of distance in the density-matrix space, which is the most obvious choice for that metric, also contributes significantly to this behavior. Even drastic interventions in this definition of distance, however, were unable to qualitatively change the POD algorithm's behavior (and resulting failing filter) to match the physical system's behavior.

At the root of concerns about asymmetry and distance is the fundamental limitation of Proper Orthogonal Decomposition: it only captures directions away from a mean state, defines a linear space, and fundamentally relates only to small perturbations away from the mean. Large dynamics away from the mean and nonlinear underlying structures of the dynamics (such as multiple stable states) all challenge POD in precisely this weak spot. Unfortunately, it is often precisely this behavior which we wish to capture or control.

Chapter 4

Dynamics on Nonlinear Manifolds

Proper Orthogonal Decomposition generates linear spaces onto which we may project system dynamics. However, the underlying space in which dynamics actually occur may not be linear. Instead, we may wish to find a more general manifold on which the dynamics occurs. To find such a manifold from a quantum trajectory, we turn to results from the field of unsupervised manifold learning. “Unsupervised” in this context means that the process of finding and defining a manifold is algorithmic, although as we will see there is still some room for researcher intervention and intuition (e.g., determining the dimension of the manifold for which the algorithm is searching, and how many neighbors to include in the definition of a point’s “neighborhood”). There have been several manifold learning algorithms proposed in recent years in the machine learning community and elsewhere (including [32, 35, 36, 37, 38, 39, 40] and references therein). These manifold learning algorithms share a focus on locally linear manifolds, which are then stitched together into a global manifold definition. As a result, the manifold ends up being defined locally, with a different function at each point mapping from the full space onto the manifold.

Of the algorithms we considered, we find the Local Tangent Space Alignment algorithm from Zhang and Zha [32] to be the most satisfying. Most prominently, its definition of the tangent space at each point has more rigorous differential geometry treatment. This is intimately related to the fact that the algorithm starts by doing what is effectively Proper Orthogonal Decomposition at each point to define the optimal linear subspace. In fact, if we were to define the full sample as “nearest neighbors” of every point, the LTSA algorithm would reduce to POD.

4.1 Local Tangent Space Alignment

The Local Tangent Space Alignment algorithm for manifold learning was introduced by Zhang and Zha [32], and in this section we will briefly review that algorithm, clearly paralleling their derivation (but choosing only to describe the form of the algorithm I use in analyzing our cavity QED dynamics). The input for the algorithm is a set of points in an m -dimensional space, which we believe may fall

on or near a d -dimensional (likely nonlinear) manifold. Noise may make these points not fall exactly on the lower-dimensional manifold, but we assume that structureless noise will average out, allowing us to extract the underlying manifold.

We will construct the manifold by aligning the tangent spaces (\mathcal{T}_τ) of each point (τ). If we had a function $f : \Omega \subset \mathbb{R}^d \rightarrow \mathbb{R}^m$ that mapped from the d -dimensional manifold Ω to the full space, we could use its Jacobian $J_f(\tau)$ at each point to find the tangent space. If Q_τ is an orthonormal basis of \mathcal{T}_τ , $\bar{\tau}$ is a neighbor of τ , and $\theta_\tau^*(\bar{\tau})$ is the local coordinate of $\bar{\tau}$ in the basis Q_τ , then

$$J_f(\tau)(\bar{\tau} - \tau) = Q_\tau \theta_\tau^*(\bar{\tau}). \quad (4.1)$$

We would like to reconstruct the coordinates τ from the data sampled in the m -dimensional space, without knowing (or being forced to calculate) f . The alignment process will allow us to go from the local coordinates θ_τ^* to the global coordinates τ . Denote $P_\tau = Q_\tau^T J_f(\tau)$, and then we have

$$\theta_\tau^*(\bar{\tau}) = Q_\tau^T J_f(\tau)(\bar{\tau} - \tau) \equiv P_\tau(\bar{\tau} - \tau). \quad (4.2)$$

(Recall that Q_τ is orthonormal so that $Q_\tau^T = Q_\tau^{-1}$.) Now P_τ is a local map from the global coordinates to the local. The Jacobian is a creation of the first derivatives of f , so the local coordinates θ_τ^* tell us about the linear part of $f(\bar{\tau}) - f(\tau)$. We can approximate them instead by the orthogonal projection of $f(\bar{\tau}) - f(\tau)$ onto \mathcal{T}_τ ; we denote these projected coordinates $\theta_\tau(\bar{\tau})$. P_τ should be invertible because the Jacobian should be full rank, so we can denote its inverse L_τ , and we see that (to first order)

$$\bar{\tau} - \tau \approx P_\tau^{-1} \theta_\tau(\bar{\tau}) = L_\tau \theta_\tau(\bar{\tau}). \quad (4.3)$$

So, instead of finding the function f , instead we are looking for a set of coordinates τ , each with a local affine transformation L_τ , to minimize

$$\int_\Omega \left(\int_{\Omega(\tau)} \|\bar{\tau} - \tau - L_\tau \theta_\tau(\bar{\tau})\| d\bar{\tau} / \int_{\Omega(\tau)} d\bar{\tau} \right) d\tau \quad (4.4)$$

over all possible L_τ .

Moving one step closer to the practical implementation of the algorithm, let us now consider implementing this process with a given finite set of data points $X = [x_1, \dots, x_N]$. In practice, the LTSA algorithm begins by defining a local neighborhood for each point, using some metric on the m -dimensional space. In the quantum trajectory case, I used the 2-norm distance between the vectorized versions of density matrices, themselves produced from trajectory wavefunctions. This is the same distance measure as in Proper Orthogonal Decomposition; however in this case we simply

want a collection of points in the neighborhood of each point, and the actual distances are not used. Another alternative would have been the 1-norm matrix distance, $|A - B|$, but I chose the 2-norm for consistency with POD and for calculation time (the 2-norm is fast to calculate, while $|A - B|$ requires calculating eigenvalues for every pair). Despite its poor scaling properties, I chose to use a brute-force method to generate these neighborhoods: I calculated the distances between every pair of points, and then sorted them (an order N^2 algorithm for N points). While some algorithms exist which scale better with respect to the number of points, they generally scale poorly in the dimension of the space. In our case, the density matrix space may have more than 10,000 dimensions, and we calculate all the distances between up to 2,000 points.

We would like to use the set of neighbors of each point τ to approximate the tangent space \mathcal{T}_τ at that point. In effect, we would like to do at each point what we did globally in Chapter 3: compute the best d -dimensional affine subspace approximation for the data points in the neighborhood of each point. Let us denote a matrix of the k nearest neighbors of x_i (including x_i) as $X_i = [x_{i_1}, \dots, x_{i_k}]$. As in our discussion of Proper Orthogonal Decomposition, we could calculate the basis of this optimal subspace by calculating the eigenvectors corresponding to the largest d eigenvalues of the correlation matrix

$$(X_i - \bar{x}_i e^T) (X_i - \bar{x}_i e^T)^T \quad (4.5)$$

where \bar{x}_i is the mean of the neighbors of x_i (recall that the optimal affine subspace will go through the mean), and e is a vector of length k and all ones (to make the dimensions work out). What we would like, instead of the definition of the basis vectors of the subspace, however, are the local coordinates of the points in that space, with that basis. These are calculated similarly, from the leading d eigenvectors of

$$(X_i - \bar{x}_i e^T)^T (X_i - \bar{x}_i e^T) \quad (4.6)$$

(The leading eigenvector, of length k , gives the coordinate of each of the k points in the direction of the leading basis vector from Eqn. (4.5), and so-on for each eigenvector.)

These local coordinates in the basis for each point, our $\theta_j^{(i)}$ s ($j = 1, \dots, k$, $i = 1, \dots, N$), tell us about the local geometry. We would like global coordinates, τ_i , which respect that geometry:

$$\tau_{i_j} = \bar{\tau}_i + L_i \theta_j^{(i)} + \epsilon_j^{(i)} \quad (4.7)$$

where $\bar{\tau}_i$ is the mean of the k τ_{i_j} s, L_i is a local affine transformation matrix (to be determined), and $\epsilon_j^{(i)}$ is the local error. If we combine these k equations into one matrix equation, and denote the local coordinates Θ_i , $T_i = [\tau_{i_1}, \dots, \tau_{i_k}]$, and $E_i = [\epsilon_1^{(i)}, \dots, \epsilon_k^{(i)}]$, then we have

$$T_i = \frac{1}{k} T_i e e^T + L_i \Theta_i + E_i. \quad (4.8)$$

We would like to minimize the error term, so we should minimize

$$\sum_{i=1}^N \|E_i\|^2 = \sum_{i=1}^N \left\| T_i \left(I - \frac{1}{k} ee^T \right) - L_i \Theta_i \right\|^2. \quad (4.9)$$

This looks familiar — again we calculate eigenvectors. However, this time we would like the eigenvectors which correspond to the smallest nonzero eigenvalues. First, though, we need to expand from the single point to consider all of the points making the full manifold. Let $T = [\tau_1, \dots, \tau_N]$ and S_i be the selection matrix so that $TS_i = T_i$. S_i can, of course, be calculated from our global distance matrix, and does not require T (which we need S_i to solve for). Construct $S = [S_1, \dots, S_N]$ and $W = \text{diag}(W_1, \dots, W_N)$ with

$$W_i = \left(I - \frac{1}{k} ee^T \right) (I - \Theta_i^+ \Theta_i) \quad (4.10)$$

where Θ_i^+ is the Moore-Penrose generalized inverse of Θ_i . Then the overall reconstruction error is given by

$$\sum_{i=1}^N \|E_i\|^2 = \|TSW\|^2. \quad (4.11)$$

This could of course be minimized by $T = 0$, so we require that $TT^T = I$. This allows us to construct the correlation matrix

$$B \equiv SWW^T S^T. \quad (4.12)$$

The vector of all ones (e) is an eigenvector of B with eigenvalue zero; the optimal T is given by the d eigenvectors with the 2nd to $d + 1$ st smallest eigenvalues.

From a numerical standpoint, what we need, then, is the matrix B . Rather than explicitly calculating with the selection matrix S , we simply calculate WW^T at each point and sum them up. Define g_1, \dots, g_d to be the unit eigenvectors corresponding to the d largest eigenvalues of Eqn. (4.6). Define $G_i = [e/\sqrt{k}, g_1, \dots, g_d]$, and one can then verify that

$$W_i = I - G_i G_i^T. \quad (4.13)$$

Now see that $W_i W_i^T = W_i$, so that B may be computed by locally summing $I - G_i G_i^T$. This is in fact what we do when solving the algorithm numerically: at each point, we find G_i , and we then sum $I - G_i G_i^T$ over the full set of points to build the full B . Computing the eigenvectors of B with the d smallest nonzero eigenvalues gives us the global coordinates of each point in our new basis, resulting in a point-by-point definition of the manifold $\Omega \subset \mathbb{R}^d$. Zhang and Zha [32] provide an error analysis of the algorithm, and provide bounds for the reconstruction errors. For our purposes, it suffices that the errors are proven to be bounded, and can be quite small.

It should now be clear that LTSA effectively reduces to Proper Orthogonal Decomposition if we define the “neighborhood” of each point to include all the points. The additional degree of freedom provided by k , the number of neighbors to use, makes the LTSA algorithm more flexible, but also less exact. The same data set, processed with different values of k , will result in different manifold definitions. This makes it difficult to define a “definitive” manifold corresponding to a set of data. However, if we run the LTSA algorithm with a range of values of k (and assume that the dimensionality d is appropriate to the data), there should be a range of k over which the general structure of the manifold does not change. We assume that this structure is the “correct” result from the algorithm, and analyze one such result as representative of that data set. If the data set is well sampled, and the dimension d chosen is reasonable, we expect that a similar manifold structure and topology would be determined from any such data set sampled from the same dynamical system.

Unlike Proper Orthogonal Decomposition, the LTSA algorithm requires fixing the value of d prior to running the algorithm. In POD, one may simply choose the dimensionality of the calculated subspace at the end, by choosing the number of eigenstates/basis vectors to use. In LTSA, however, we calculate the d -dimensional space at each point, and then integrate them together (the form of G_i depends on d). While the costly distance calculation may be reused for different d s, the bulk of the algorithm must be run separately for each d . In practice, we do not know d in advance, and must run the algorithm multiple times, to see at what point d is sufficient. For example, take a case for which the inherent dimension of the manifold is 2. When run with $d = 3$, the resulting manifold should show no interesting structure in the third dimension: it should look like a flat 2-dimensional surface suspended in 3 dimensions. This is straightforward to see on screen for 2 and 3 dimensions; for higher dimensions one has to look at the variation across the sample in the d^{th} dimension, and make a decision as to whether it is simply noise from the input data set. In the end, we are simply lucky that the underlying structures for the dynamics in the cavity QED cases we consider seem to fit well in few dimensions, which allows us to proceed to build approximations to those dynamical systems which reside on or close to these manifolds.

4.2 Deriving dynamical systems on LTSA manifolds

Once we have a manifold determined from a set of trajectories, processed by the LTSA algorithm, we would like to derive a set of equations for the dynamics of reduced-order systems which live on that manifold. Unfortunately, because of the local nature of the LTSA algorithm, we do not have a global function which we can use to map the dynamics of the system of equations which generated the trajectories into dynamics on the lower-dimensional manifold. In particular, we cannot use the filter projection technique discussed in Chapter 2 without a functional form for the maps between the coordinates on the d -dimensional manifold and the coordinates in the m -dimensional space of

trajectories.

One option would be to find and utilize the local maps between the tangent spaces of each point and the larger space. This would be highly accurate (not much of an approximation). However, it would be impractical: as we calculate a new trajectory, we would need to find the nearest point from the originating trajectory at each time step, project the dynamical equations into its particular space, advance the point in local coordinates, revert to global coordinates, and repeat.

Another option, this time approximate, would be to attempt to utilize a least-squares fitting algorithm to find a suitable map from the d -dimensional manifold to the m -dimensional originating space. In the case of cavity QED, however, m is very large. For example, $m = 14,400$ for a 120×120 density matrix, while (as we will see below), d may be less than 5. More than m fit parameters would need to be calculated, requiring more than m samples from the trajectory. In addition, there is no clear way to decide *a priori* what functional form to assume for the fitting function, meaning either a very high-dimensioned guess with a linear form, or building out polynomials (with the requisite larger number of fit parameters).

The process I have chosen instead is to relax the requirement to reproduce the full density matrix from the global manifold coordinates (and *vice versa*), and instead focus only on the expectation values of operators. Here we are effectively assuming that the simple operators (those we have the possibility to measure) carry the information we care about regarding the system, and hoping that the manifold can be easily fit by a simple polynomials of these operators.

Carrying this plan forward, let us assume we have equations mapping expectation values to manifold coordinates, similar to this:

$$\tau_1 = v_1 \langle x \rangle + v_2 \langle y \rangle + v_3 \langle \sigma_x \rangle + \dots + v_k \langle xy\sigma_x \rangle \langle x \rangle + \dots, \quad (4.14)$$

where v_j are fit parameters, and we want an equation for motion for τ_1 . That is, we would like an expression for $d\tau_1$ (and similar for $d\tau_2$ to $d\tau_d$).

In order to generate these equations, we need a general prescription for $d\langle \mathcal{O} \rangle$.

$$d\langle \mathcal{O} \rangle = d(\text{Tr}(\mathcal{O}\rho)) \quad (4.15)$$

$$= \text{Tr}(d(\mathcal{O}\rho)) \quad (4.16)$$

$$= \text{Tr}((d\mathcal{O})\rho + \mathcal{O}(d\rho)). \quad (4.17)$$

In the Schrödinger picture, $d\mathcal{O} = 0$ while $d\rho \neq 0$, so we really just need to know

$$d\langle \mathcal{O} \rangle = \text{Tr}(\mathcal{O}(d\rho)). \quad (4.18)$$

For our Cavity QED system, we substitute the full Stochastic Master Equation, and see that we

will need to calculate

$$\begin{aligned}
d\langle \mathcal{O} \rangle &= \text{Tr}(-i\mathcal{O}[H, \rho]dt + \kappa\mathcal{O}(2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a)dt \\
&\quad + \gamma\mathcal{O}(2\sigma\rho\sigma^\dagger - \sigma^\dagger\sigma\rho - \rho\sigma^\dagger\sigma)dt \\
&\quad + i\sqrt{2\kappa}\mathcal{O}(\rho a^\dagger - a\rho - \text{Tr}[\rho(a^\dagger - a)]\rho)dW)
\end{aligned} \tag{4.19}$$

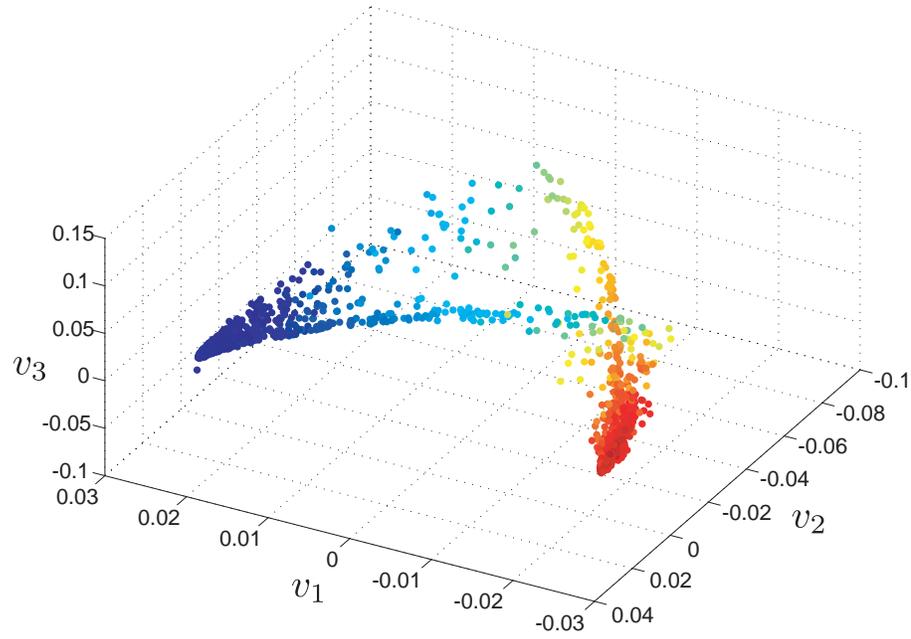
$$\begin{aligned}
&= -i\langle \mathcal{O}H \rangle dt + i\langle H\mathcal{O} \rangle dt + 2\kappa\langle a^\dagger\mathcal{O}a \rangle dt - \kappa\langle \mathcal{O}a^\dagger a \rangle dt \\
&\quad - \kappa\langle a^\dagger a\mathcal{O} \rangle dt + 2\gamma\langle \sigma^\dagger\mathcal{O}\sigma \rangle dt - \gamma\langle \mathcal{O}\sigma^\dagger\sigma \rangle dt - \gamma\langle \sigma^\dagger\sigma\mathcal{O} \rangle dt \\
&\quad + i\sqrt{2\kappa}\langle a\mathcal{O} \rangle dW + \sqrt{2\kappa}\langle \mathcal{O}a^\dagger \rangle dW - \sqrt{2\kappa}\langle a^\dagger - a \rangle \langle \mathcal{O} \rangle dW.
\end{aligned} \tag{4.20}$$

This plan depends on finding understandable, finite sets of expectation values from which to determine the manifold coordinates. I found such a set of expectation values for the simpler case of phase bistability, but did not for absorptive bistability.

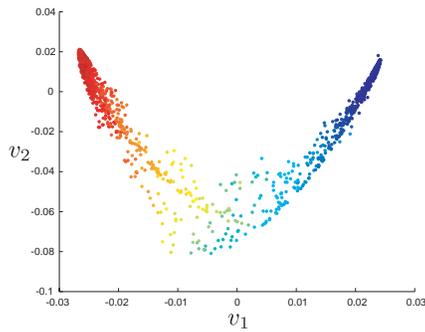
4.3 Phase bistability: Dynamics on three dimensional manifolds from LTSA

The full quantum dynamics of cavity QED, our model open quantum system, take place in an infinite-dimensional space, which I have approximated in this research by a space of density matrices with tens of thousands of elements. Quantum trajectory simulations preserve pure states, so the space in which they take place is much smaller. However, if we keep the first 150 Fock states, as I have in my simulations, the wavefunction still resides in a space with hundreds of dimensions. Thankfully, the dynamics of our canonical phase bistable system use only a small fraction of those degrees of freedom, which we can identify using the LTSA algorithm. Unlike Proper Orthogonal Decomposition, the full algorithm must be run for each potential subspace dimensionality (that is, run separately for $d = 1$, $d = 2$, etc). I ran the LTSA algorithm on a selection of 2000 samples from the same long trajectory analyzed by POD in Section 3.3. $k = 60$ nearest neighbors proved to be in a broad range of k values with similar manifold topology.

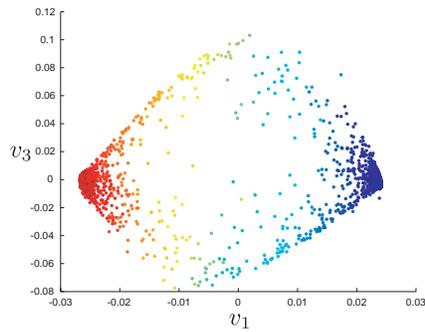
For $d = 1$, the manifold consists of a line segment with samples from one stable point at one end, and the other stable point at the other. It is almost monotonic in $\langle y \rangle$. Expanding to two dimensions ($d = 2$), the manifold takes a roughly parabolic shape. However, at the base of the parabola (in the transition between stable points), the samples align in two bands, which correspond to different $\langle \sigma_y \rangle$. Recall that physically the jumps between phase states correspond to atomic spontaneous emissions events: the atom switches before the field, with the field switching slower because of the cavity. The two bands at the base of the parabola, therefore, correspond to upward-going and downward-going transitions.



(a)



(b)



(c)

Figure 4.1: Local Tangent Space Alignment 3-dimensional manifold from a long quantum trajectory in the phase bistable regime. LTSA algorithm run on 2,000 samples from the trajectory, with $k = 60$ nearest neighbors. Scatter points are colored according to the value of $\langle y \rangle$, the expectation value of the phase quadrature of the cavity field. Part (a) shows a 3D plot; (b) shows just v_2 plotted against v_1 ; (c) shows v_3 versus v_1 . In (c), the upper and lower bands correspond to upward-going and downward-going transitions between the two stable states. Trajectory parameters: $\Theta = 0$, $\Delta = 0$, $\kappa = 4$, $\gamma = 2$, $g_0 = 12$, $\mathcal{E} = 23.57$, $t_{final} = 5000$.

Id	x	y	σ_x	σ_y
σ_z	x^2	$xy + yx$	$x\sigma_x$	$x\sigma_y$
$x\sigma_z$	y^2	$y\sigma_x$	$y\sigma_y$	$y\sigma_z$

Table 4.1: The 15 unique first- and second-order Hermitian operators constructed from the Identity and the five operators from the Maxwell Bloch equations.

The continued structure in the second dimension of the LTSA manifold prompts us to expand to $d = 3$, which produces the manifold show in Figure 4.1. The first two dimensions remain qualitatively similar to their appearance at lower dimensions, while the third dimension shows a split between the two transition bands described above. Extending to $d = 4$, however, does not provide additional insight: this fourth coordinate is apparently randomly scattered with respect to the first three, and shows no structure related to system expectation values. I therefore conclude that the underlying manifold for this regime is very likely three dimensional, and continue with further analysis of only the $d = 3$ manifold.

As mentioned in Section 4.2, I attempted to fit the three dimensions which define this manifold with polynomial combinations of operator expectation values. In particular, I built the set of all independent first and second order Hermitian operators which can be created from combinations of the five operators familiar from the Maxwell-Bloch equations: x and y the field quadratures, and σ_x , σ_y , and σ_z the atomic operators. This set of 15 independent Hermitian operators is shown in Table 4.1. If we call these operators in general \mathcal{O}_i , then I attempted to fit each of the three global manifold coordinates $v_i, i = 1, 2, 3$ with quadratic combinations

$$v_i = \sum_{j=1}^{15} \sum_{k=j}^{15} x_{jk} \langle \mathcal{O}_j \rangle \langle \mathcal{O}_k \rangle. \quad (4.21)$$

Note that one of the operators is the Identity, so this includes linear terms.

As evident from Figure. 4.1, the leading dimension, v_1 , of this LTSA manifold has a very strongly monotonic relationship with the phase quadrature of the cavity field, $\langle y \rangle$. The two stable points reside at either end of the arc, with points from transitions in between. The roughly parabolic shape of the arc in Figure 4.1b leads us to guess that $\langle y^2 \rangle$ or $\langle y \rangle^2$ will play a significant role in approximations of v_2 . However, rather than simply guess, I used a “basis pursuit” technique (also known as “ l_1 -form regularization”) to algorithmically identify the leading contributors to each dimension [41]. The CVX convex optimization library for Matlab [42] makes this technique straightforward.

A normal least-squares fit involves minimizing the l_2 norm of the difference between the fit and the data, $\|Ax - v\|_2$, where $A \in \mathbb{R}^{m \times n}$ (the columns of A are the values of the basis functions we’re using to fit the data, called the regressors), $v \in \mathbb{R}^m$ are the data, and $x \in \mathbb{R}^n$ are the fit coefficients (the variable we’re solving for). *Basis pursuit* minimizes the sum of this l_2 norm and the l_1 norm of the coefficients of the fit. In general, minimizing an l_1 norm puts pressure on the

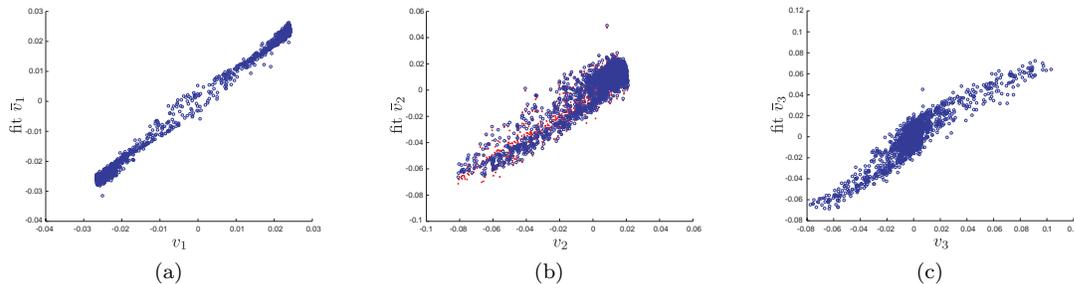


Figure 4.2: Scatter plots of 2,000 data points showing the fit between Local Tangent Space Alignment coordinates v_i and fits to those coordinates using only a few expectation values. (a) v_1 fit by $\langle y \rangle$; (b) v_2 fit by $\langle y^2 \rangle$ (blue circles) and by $\langle y^2 \rangle$ and $\langle y \sigma_y \rangle$ (red crosses); (c) v_3 fit by $\langle x \sigma_y \rangle$, $\langle y \rangle$, $\langle y \sigma_y \rangle$, and $\langle x \sigma_y \rangle \langle y^2 \rangle$.

fit to have coefficients which are exactly zero. This allows us to eliminate some coefficients for variables which duplicate the contributions of other variables by selecting the single or few variables which best fit the data. The balance between the l_2 and l_1 norm terms is set by the parameter γ in minimizing $(\|Ax - v\|_2 + \gamma\|x\|_1)$. I weighted the l_1 term by the root mean square of each corresponding expectation value; this controls for the fact that the coefficient of a much larger-valued operator will necessarily be smaller, meaning a naked basis pursuit process will eliminate it even if it is a better fit. I used the basis pursuit process to narrow the number of variables by increasing γ until the number of surviving variables fell to a reasonable number (between one and five) while still allowing a good fit to the data. The presence of the l_1 term can skew the results for the optimal values of the least-squares parameters, so after settling on a small set of variables to include, I did a simple l_2 norm (least-squares) fit with only that parameter set.

Applying the basis pursuit process to the leading dimension of the 3-dimensional phase bistability manifold, v_1 , the number of included operators shrinks rapidly as γ is increased. 11 operators survive at $\gamma = 5$, but $\langle y \rangle$ is already clearly dominant. At $\gamma = 14$, only $\langle y \rangle$ remains. Figure 4.2a shows the almost-identity relationship between the values of v_1 and the fit

$$\hat{v}_1 = -0.0011 - 0.0178 \langle y \rangle. \quad (4.22)$$

The second dimension, as predicted, is related to $\langle y^2 \rangle$. At $\gamma = 5$, 14 operators remain, dominated by $\langle y \sigma_y \rangle$. At $\gamma = 20$, there are 6 terms with non-zero coefficients; the leading three are $\langle y^2 \rangle$, $\langle y \sigma_y \rangle$, and $\langle y^2 \rangle \langle y \sigma_y \rangle$. These three all carry roughly the same information ($\langle y \rangle$ and $\langle \sigma_y \rangle$ are largely anti-correlated). At $\gamma = 30$, only $\langle y^2 \rangle$ and $\langle y \sigma_y \rangle$ survive. For $\gamma \geq 38$, only $\langle y^2 \rangle$ survives. The best fit with just $\langle y^2 \rangle$ is shown in the blue circles in Figure 4.2b; the functional form is

$$\hat{v}_2 = -0.0767 + 0.0402 \langle y^2 \rangle. \quad (4.23)$$

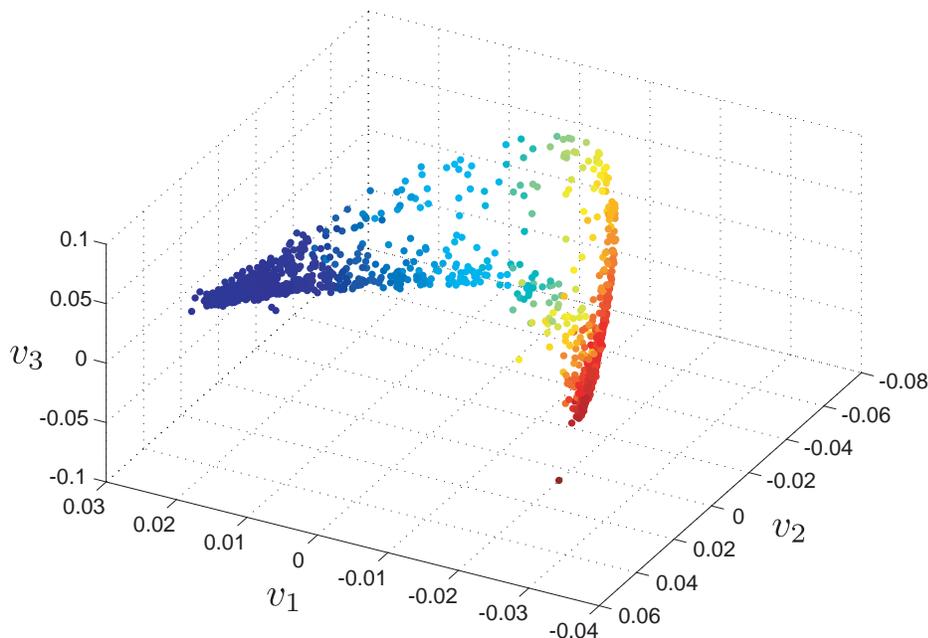


Figure 4.3: Three dimensional fit to the LTSA phase bistability manifold, from Eqns. (4.22), (4.23), and (4.25). Compare with Figure 4.1a to evaluate the fit.

If we allow $\langle y\sigma_y \rangle$ to survive as well (red crosses in Figure 4.2b), the best fit is

$$\tilde{v}_2 = -0.0811 + 0.0284 \langle y^2 \rangle - 0.0244 \langle y\sigma_y \rangle. \quad (4.24)$$

In contrast to the leading two dimensions, the third is not a clear function of any one or two parameters. Basis pursuit narrows the search down to $\langle x\sigma_y \rangle$, $\langle y \rangle \langle y\sigma_y \rangle$ and $\langle x\sigma_y \rangle \langle y^2 \rangle$. The best fit is shown in comparison with the data in Figure 4.2c, and its functional form is

$$\hat{v}_3 = 8.43 \times 10^{-4} - 9.52 \times 10^{-3} \langle x\sigma_y \rangle + 8.26 \times 10^{-2} \langle y \rangle \langle y\sigma_y \rangle - 9.314 \times 10^{-3} \langle x\sigma_y \rangle \langle y^2 \rangle. \quad (4.25)$$

Figure 4.3 shows the 3-dimensional manifold from the fit; compare with Figure 4.1a to see that just a few expectation values carry almost all of the information necessary to reproduce the manifold originally created by analyzing the full density matrices.

4.3.1 A new set of “Maxwell-Bloch” equations

Based on simulations and the LTSA algorithm, I have determined that the first two dimensions of the manifold on which phase bistability dynamics take place can be well approximated with only $\langle y \rangle$ and $\langle y^2 \rangle$. We would like to build analytic expressions for system dynamics on this manifold. In

order to build a complete set of equations for the dynamics of expectation values, we will need to add the three directions on the atomic Bloch sphere as our other observables. We can assume that, because y and y^2 do such a good job, alone, with the leading two dimensions of the phase bistability dynamics, we will be able to make reasonable assumptions to remove other terms and derive a closed set of equations, and in fact we can.

Let us now derive the equations of motion for five expectation values, $\langle y \rangle$, $\langle y^2 \rangle$, $\langle \sigma_x \rangle$, $\langle \sigma_y \rangle$, and $\langle \sigma_z \rangle$. This process looks very much like the derivation of the Maxwell-Bloch equations, except that we're going to make a few assumptions (based on the simulation results) along the way to simplify things.

First, let us calculate $\langle y \rangle$, noting that y commutes with the \mathcal{E} term in the Hamiltonian, as well as all σ operators. Also, we have set the detuning parameters to zero, because that is what they are for phase bistable behavior. The commutator of y and the Hamiltonian is

$$\begin{aligned}
[y, H] &= ig_0[y, a^\dagger \sigma_- - a \sigma_+] \\
&= ig_0 (ya^\dagger \sigma_- - ya \sigma_+ - a^\dagger \sigma_- y + a \sigma_+ y) \\
&= ig_0 ((ya^\dagger - a^\dagger y) \sigma_- + (ay - ya) \sigma_+) \\
&= -ig_0 \left(\frac{i}{2} \sigma_- + \frac{i}{2} \sigma_+ \right) \\
&= \frac{-ig_0}{2} \sigma_y.
\end{aligned} \tag{4.26}$$

Therefore, the equation for the motion of $\langle y \rangle$ is

$$\begin{aligned}
d\langle y \rangle &= -i\langle [y, H] \rangle dt + \kappa \langle 2a^\dagger ya - a^\dagger ay - ya^\dagger a \rangle dt \\
&\quad + \sqrt{2\kappa} \langle a^\dagger y - ya - \langle a^\dagger - a \rangle y \rangle dW \\
&= \frac{-g_0}{2} \langle \sigma_y \rangle dt - \kappa \langle y \rangle dt + \sqrt{8\kappa} \left(\langle y^2 \rangle - \langle y \rangle^2 \right) dW.
\end{aligned} \tag{4.27}$$

Note that we could have derived this (aside from the dW term) from the $\langle \dot{a} \rangle$ equation in the standard Maxwell Bloch equations. The dW always takes a standard form related to the difference between $\langle \mathcal{O}y \rangle$ and $\langle \mathcal{O} \rangle \langle y \rangle$, so we can avoid calculating the equations of motion for $\langle \sigma_{x,y,z} \rangle$, and just state them here:

$$d\langle \sigma_x \rangle = 2g_0 \langle x \sigma_z \rangle dt - \gamma \langle \sigma_x \rangle dt + \sqrt{8\kappa} (\langle y \sigma_x \rangle - \langle y \rangle \langle \sigma_x \rangle) dW \tag{4.28}$$

$$d\langle \sigma_y \rangle = -2g_0 \langle y \sigma_z \rangle dt - \gamma \langle \sigma_y \rangle dt + \sqrt{8\kappa} (\langle y \sigma_y \rangle - \langle y \rangle \langle \sigma_y \rangle) dW \tag{4.29}$$

$$\begin{aligned}
d\langle \sigma_z \rangle &= -g_0 (\langle x \sigma_x \rangle - \langle y \sigma_y \rangle) dt - 2\gamma (\langle \sigma_z \rangle + 1) dt \\
&\quad + \sqrt{8\kappa} (\langle y \sigma_z \rangle - \langle y \rangle \langle \sigma_z \rangle) dW.
\end{aligned} \tag{4.30}$$

The $\langle y^2 \rangle$ equation is new, and cannot be simply written down based on the Maxwell Bloch equations. y^2 commutes with the same terms in the Hamiltonian that y does, simplifying that commutation calculation.

$$\begin{aligned}
[y, H] &= ig_0[y^2, a^\dagger \sigma_- - a \sigma_+] \\
&= ig_0(y^2 a^\dagger \sigma_- - y^2 a \sigma_+ - a^\dagger \sigma_- y^2 + a \sigma_+ y^2) \\
&= ig_0((y^2 a^\dagger - a^\dagger y^2) \sigma_- + (ay^2 - y^2 a) \sigma_+) \\
&= -g_0(y \sigma_+ - y \sigma_-) \\
&= -ig_0 y \sigma_y
\end{aligned} \tag{4.31}$$

where we have used the commutation relations between y^2 and a and a^\dagger :

$$[y^2, a] = yya - ayy = y(ay - \frac{i}{2}) - (ya + \frac{i}{2})y = -iy \tag{4.32}$$

$$[y^2, a^\dagger] = yya^\dagger - a^\dagger yy = y(a^\dagger y - \frac{i}{2}) - (ya^\dagger + \frac{i}{2})y = -iy. \tag{4.33}$$

Looking now to the full differential equation, and plugging in the result for the commutator, we have

$$\begin{aligned}
d\langle y^2 \rangle &= -i\langle [y, H] \rangle dt + \kappa \langle 2a^\dagger y^2 a - a^\dagger ay^2 - y^2 a^\dagger a \rangle dt \\
&\quad + \sqrt{2\kappa} \langle a^\dagger y^2 - y^2 a - \langle a^\dagger - a \rangle y^2 \rangle dW \\
&= -g_0 \langle y \sigma_y \rangle dt + \kappa \langle a^\dagger (ay^2 - iy) - a^\dagger ay^2 + (y^2 a^\dagger + iy) a - y^2 a^\dagger a \rangle dt \\
&\quad + \sqrt{8\kappa} \langle y^3 - \langle y \rangle y^2 \rangle dW \\
&= -g_0 \langle y \sigma_y \rangle dt - i\kappa \langle a^\dagger y - ya \rangle dt + \sqrt{8\kappa} (\langle y^3 \rangle - \langle y \rangle \langle y^2 \rangle) dW \\
&= -g_0 \langle y \sigma_y \rangle dt - i\kappa \langle a^\dagger y - (ay - i/2) \rangle dt + \sqrt{8\kappa} (\langle y^3 \rangle - \langle y \rangle \langle y^2 \rangle) dW \\
&= -g_0 \langle y \sigma_y \rangle dt - 2\kappa \langle y^2 \rangle dt + \kappa/2 dt + \sqrt{8\kappa} (\langle y^3 \rangle - \langle y \rangle \langle y^2 \rangle) dW.
\end{aligned} \tag{4.34}$$

Collecting our five equations together, we now have:

$$d\langle y \rangle = \frac{-g_0}{2} \langle \sigma_y \rangle dt - \kappa \langle y \rangle dt + \sqrt{8\kappa} (\langle y^2 \rangle - \langle y \rangle^2) dW \tag{4.35}$$

$$d\langle y^2 \rangle = -g_0 \langle y \sigma_y \rangle dt - 2\kappa \langle y^2 \rangle dt + \kappa/2 dt + \sqrt{8\kappa} (\langle y^3 \rangle - \langle y \rangle \langle y^2 \rangle) dW \tag{4.36}$$

$$d\langle \sigma_x \rangle = 2g_0 \langle x \sigma_z \rangle dt - \gamma \langle \sigma_x \rangle dt + \sqrt{8\kappa} (\langle y \sigma_x \rangle - \langle y \rangle \langle \sigma_x \rangle) dW \tag{4.37}$$

$$d\langle \sigma_y \rangle = -2g_0 \langle y \sigma_z \rangle dt - \gamma \langle \sigma_y \rangle dt + \sqrt{8\kappa} (\langle y \sigma_y \rangle - \langle y \rangle \langle \sigma_y \rangle) dW \tag{4.38}$$

$$\begin{aligned}
d\langle \sigma_z \rangle &= -g_0 (\langle x \sigma_x \rangle - \langle y \sigma_y \rangle) dt - 2\gamma (\langle \sigma_z \rangle + 1) dt \\
&\quad + \sqrt{8\kappa} (\langle y \sigma_z \rangle - \langle y \rangle \langle \sigma_z \rangle) dW.
\end{aligned} \tag{4.39}$$

Note that this set of equations is not closed: there are many terms on the right hand side which have no corresponding dynamical equation. This is where approximation and calculations from the simulations come in. I fit the expectation values of these operators using basis pursuit followed by a least-squares fit:

$$\begin{aligned}
\langle y\sigma_x \rangle &\approx A_0 + A_1 \langle y \rangle \langle \sigma_x \rangle \\
\langle y\sigma_y \rangle &\approx B_0 + B_1 \langle y \rangle \langle \sigma_y \rangle \\
\langle y\sigma_z \rangle &\approx C_0 + C_1 \langle y \rangle \langle \sigma_z \rangle \\
\langle x\sigma_x \rangle &\approx D_0 + D_1 \langle x \rangle \langle \sigma_x \rangle = D_0 + D_1(E_0 + E_1 \langle y^2 \rangle) \langle \sigma_x \rangle \\
&\quad H_0 + H_1 \langle \sigma_x \rangle + H_2 \langle y^2 \rangle \langle \sigma_x \rangle \\
\langle x\sigma_z \rangle &\approx F_0 + F_1 \langle x \rangle \langle \sigma_z \rangle = F_0 + F_1(E_0 + E_1 \langle y^2 \rangle) \langle \sigma_z \rangle \\
&\quad = J_0 + J_1 \langle \sigma_z \rangle + J_2 \langle y^2 \rangle \langle \sigma_z \rangle \\
\langle y^3 \rangle &\approx G_0 + G_1 \langle y \rangle + G_2 \langle y \rangle \langle y^2 \rangle.
\end{aligned} \tag{4.40}$$

All the fits which support these approximations are quite good (although there is significant scatter on some, there is no clear break from a linear response). The best fit values for these constants are:

$$\begin{aligned}
A_0 &= 0.0059 & B_0 &= -0.4392 & C_0 &= 0.0003 \\
A_1 &= 1.0060 & B_1 &= 0.6706 & C_1 &= 0.8751 \\
H_0 &= 0.0217 & J_0 &= -0.0468 & G_0 &= 0.0071 \\
H_1 &= 5.8911 & J_1 &= 5.8120 & G_1 &= 0.8526 \\
H_2 &= -0.1602 & J_2 &= -0.1213 & G_2 &= 0.8487
\end{aligned}$$

Renaming variables to simplify the display and emphasize that we're now dealing with a simple set of classical stochastic differential equations, we now have

$$\begin{aligned}
dv_1 &= \frac{-g_0}{2} v_4 dt - \kappa v_1 dt + \sqrt{8\kappa} (v_2 - v_1^2) dW \\
dv_2 &= -g_0 (B_0 + B_1 v_1 v_4) dt - 2\kappa v_2 dt + \kappa/2 dt \\
&\quad + \sqrt{8\kappa} (G_0 + G_1 v_1 + (G_2 - 1) v_1 v_2) dW \\
dv_3 &= 2g_0 (J_0 + J_1 v_5 + J_2 v_2 v_5) dt - \gamma v_3 dt \\
&\quad + \sqrt{8\kappa} (A_0 + (A_1 - 1) v_1 v_3) dW \\
dv_4 &= -2g_0 (C_0 + C_1 v_1 v_5) dt - \gamma v_4 dt + \sqrt{8\kappa} (B_0 + (B_1 - 1) v_1 v_4) dW \\
dv_5 &= g_0 ((B_0 + B_1 v_1 v_4) - (H_0 + H_1 v_3 + H_2 v_2 v_3)) dt - 2\gamma (v_5 + 1) dt \\
&\quad + \sqrt{8\kappa} (C_0 + (C_1 - 1) v_1 v_5) dW.
\end{aligned} \tag{4.41}$$

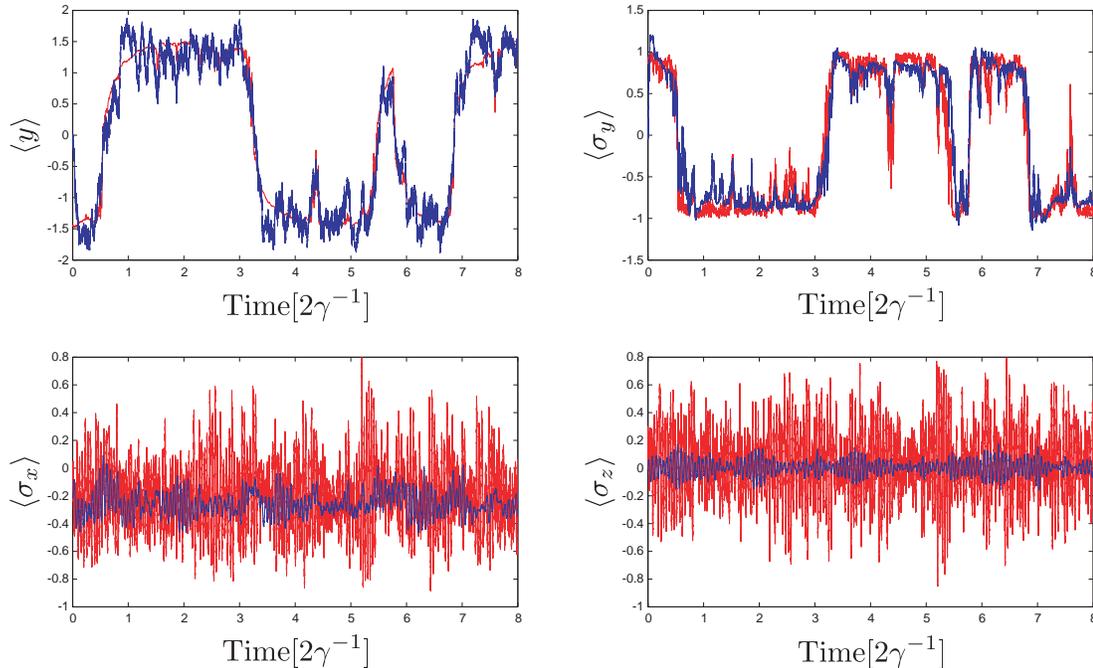


Figure 4.4: Excerpt traces for the expectation values of $\langle y \rangle$, $\langle \sigma_y \rangle$, $\langle \sigma_x \rangle$, and $\langle \sigma_z \rangle$ for a quantum trajectory simulation (red) and the filter in Eqn. (4.41) (blue). The output of the filter in Eqn. (4.42) would be virtually indistinguishable for $\langle y \rangle$ and $\langle \sigma_y \rangle$; that filter has no role for $\langle \sigma_x \rangle$ and $\langle \sigma_z \rangle$.

Simulations of these equations with white noise input for dW show that v_1 fluctuates in the range of -2 to 2, with possible stable points at near ± 1.5 (close to where we would expect the stable points of the quantum system to be). The distribution of v_1 is bimodal, with peaks at ± 1.5 .

These equations may be used as a filter as well. Unfortunately, when run as a filter, with dW equal to the innovation from a quantum trajectory's homodyne photocurrent, these equations do not tend to be very stable. Sections of the photocurrent appear able to drive the filter into unphysical states more quickly than they can recover. The traces shown in Figure 4.4 were selected to lie between these points of instability. A practical implementation of this filter would need to be on the lookout for these points, and reset the filter quickly. The lack of a full density matrix means that we cannot evaluate this filter on the basis of the fidelity, like we could in the case of Proper Orthogonal Decomposition. However, we may see how it compares with POD in predicting expectation values, comparing Figure 4.4 with Figure 3.3 and Figure 3.4.

Van Handel and Mabuchi [11] use a projection filter to derive a simple set of three SDEs which also act as a very good filter for the phase bistable regime. They, however, make a fundamentally different choice (driven by their desire for an analytic solution): the manifold onto which they project requires that $\langle \sigma_x \rangle$ and $\langle \sigma_z \rangle$ are both exactly 0 (they fix the atom in a state polarized along σ_y). They are able to do a good job of reproducing the behavior of the optimal filter for $\langle y \rangle$ and $\langle \sigma_y \rangle$, but are fundamentally unable to handle any situation in which the atom has components in another

direction. It is true that in simulations of the filter equations I have derived above, the values of $\langle\sigma_x\rangle$ and $\langle\sigma_z\rangle$ appear to be uncorrelated with the switching behavior. The average of $\langle\sigma_z\rangle$ is close to zero for any given realization, while the mean of $\langle\sigma_x\rangle$ is small but nonzero.

Following the example of [11], we decouple our equations from $\langle\sigma_x\rangle$ and $\langle\sigma_z\rangle$, by setting $C_0 = 0$ and $C_1 = 0$, and are left with:

$$\begin{aligned}
dv_1 &= \frac{-g_0}{2}v_4dt - \kappa v_1dt + \sqrt{8\kappa}(v_2 - v_1^2)dW \\
dv_2 &= -g_0(B_0 + B_1v_1v_4)dt - 2\kappa v_2dt + \kappa/2dt \\
&\quad + \sqrt{8\kappa}(G_0 + G_1v_1 + (G_2 - 1)v_1v_2)dW \\
dv_4 &= -\gamma v_4dt + \sqrt{8\kappa}(B_0 + (B_1 - 1)v_1v_4)dW.
\end{aligned} \tag{4.42}$$

(Strictly speaking we are approximating that $\langle y\sigma_z\rangle = 0$ rather than $\langle\sigma_z\rangle = 0$, but this is a consistent approximation.) When it comes to filtering, the output of this filter is almost indistinguishable from the output of the 5-dimensional filter above (the three expectation values differ on average by less than 3%). Decoupling from the x and z components of the atom is clearly a very good approximation. It also seems to make numerical simulation mildly more stable, but does not solve the filter instability problem. This three-dimensional filter is also consistent with the experience of the LTSA manifold generation process: three dimensions appears to be sufficient. It is unfortunate that we were forced to develop a filter only in terms of system observables, which prevents us from evaluating the full quantum fidelity of the approximation. However, the filter appears to perform very well with the observables a experiment is likely to measure and care about, and we cannot deny the value of an easy to implement, exact functional form. These ‘‘Maxwell-Bloch’’ equations are a product of this particular trajectory, and these particular parameters. This process would need to be repeated to generate any other filter for experimental parameters.

4.4 Two and three dimensional manifolds for absorptive bistability

Absorptive bistability in cavity QED is arguably a more complicated process than phase bistability. The two stable states are not symmetric, and in fact the upper state is much broader than the vacuum, with significant variation within the state (or zone). The transitions do not have a obvious model or explanation (like the phase bistable case does with atomic spontaneous emission). It is therefore somewhat of a surprise that the Local Tangent Space Alignment-generated manifold for our canonical absorptive bistability conditions is quite simple.

I generated $d = 2$, $d = 3$, $d = 5$, and $d = 8$ manifolds with $k = 40$ nearest neighbors from

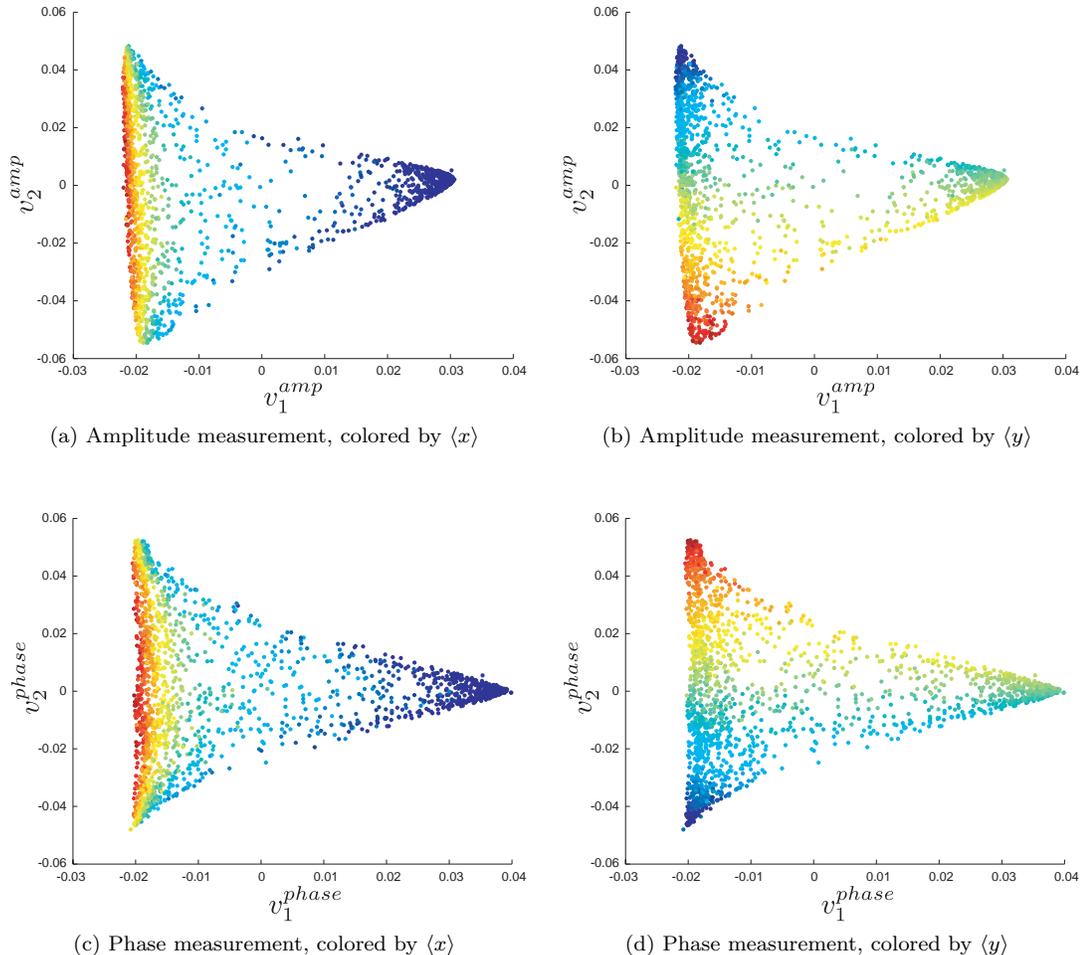


Figure 4.5: Two-dimensional Local Tangent Space Alignment-produced manifolds for absorptive bistability. (a) and (b) are from a trajectory in which the amplitude quadrature is measured by homodyne detection; (c) and (d) are from a trajectory in which the phase quadrature was measured. System parameters are our canonical set; LTSA parameters are $d = 2$ and $k = 40$, on a set of 2000 samples.

the same trajectories I used for Proper Orthogonal Decomposition, for both amplitude and phase quadrature homodyne measurement of the cavity field. The $d = 2$ manifold (Figure 4.5) shows a clear triangular-like structure, with the lower state at one point of the triangle, and the extended upper state stretching along the opposing edge. This edge corresponds roughly to variation in $\langle y \rangle$ within the upper state, and the points of highest amplitude quadrature reside in the middle of that edge. These two dimensions can be mapped one-to-one onto the phase and amplitude quadratures of the quantum trajectory points used to generate the manifold. The first dimension corresponds roughly to the amplitude quadrature variation (or perhaps more accurately the distance from the vacuum state), and the second to the phase. The eigenvalues of B for this case show that the first dimension is dominant relative to the second: the corresponding eigenvalue is roughly a factor of 7.5 smaller for amplitude measurement, and half the size for phase measurement (recall that in LTSA

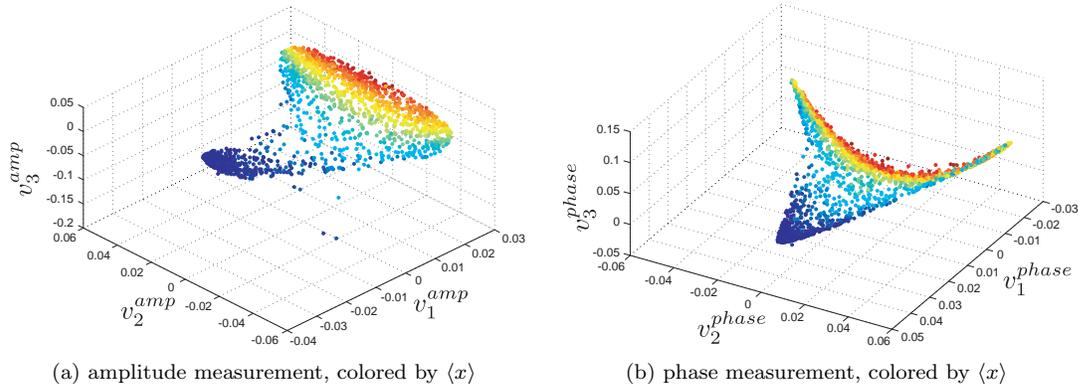


Figure 4.6: Three-dimensional Local Tangent Space Alignment-produced manifolds for absorptive bistability.

what matters is the small eigenvalue states, not the large eigenvalues as in POD). The simple atomic expectation values (σ_x , σ_y , and σ_z) do not have an obvious relation to these two dimensions.

Extending to $d = 3$, shown in Figure 4.6, we see that structure continues in this dimension, although the manifold now looks very much like a two-dimensional surface embedded in three dimensions. This may correspond to the continued growing eigenvalues from the second to the third; the eigenvalue for the third dimension is roughly 50% larger than that for the second in the fuzzier amplitude-measurement case, and over twice as large for the more crisply shaped phase quadrature manifold. The shape of the $d = 2$ manifold is very similar regardless of whether we measure the phase or amplitude quadrature (see Figure 4.5). The shape in the third dimension, however, is quite different: in the phase measurement case, the points with largest $\langle y \rangle$ curve away in the third dimension; in the amplitude measurement case, the whole upper state curves together. Perhaps this dimension corresponds to the absolute value of the measured quadrature (or possibly the square of the observable given the somewhat parabolic shape of the curve).

If we run the same process with $d = 5$, we encounter some mild structure within these added dimensions. However, the third (as before) and fourth dimensions do not show any clear relation with system observables. In the case of amplitude quadrature measurement, the fifth dimension shows a similar slight amount of structure and is monotonic with σ_y ; a similar thing happens for the sixth dimension for phase quadrature measurement. This leads me to think that the intermediate dimensions cannot be entirely overlooked, although their relation to system observables and dynamics remains somewhat of a mystery. For amplitude quadrature measurement, the eigenvalue for the fifth dimension is about 25% larger than those for the third and fourth dimensions (which are quite close together and each about 12 times larger than that for the leading dimension).

With $d = 8$, we learn that the fifth dimension was the last to show any useful structure for amplitude measurement, and the sixth for phase measurement. In both cases, the eigenvalues for the 6th, 7th, and 8th dimensions rise significantly relative to the lower dimensions, and the scatter

of points for each of the higher dimensions show that they pull out a handful of points to large displacement from the rest of the data: we have reached the threshold where structureless scatter in the data dominates.

While the values of the eigenvalues of B change as we raise d , the approximate ratios of eigenvalues appear to be quite robust to changes in d . They are slightly more dependent on k . Changing k from 40 to 60 in the amplitude measurement case, for example, changes the ratio of the leading pair of eigenvalues only from 7.5 to 8.6, and the third and fourth eigenvalues are still very close.

The manifold for absorptive bistability does not show the same clear dimension that the phase bistable case demonstrates. The underlying dynamics may require five or six dimensions (which is still a dramatic reduction compared with the size of the density matrix space), or it is even possible that different parts of the dynamics (such as the upper and lower “states”) require different numbers of dimensions. Regardless, in this regime the leading two dimensions seem to dominate, and remain reassuringly independent of the measurement used, implying that they reflect behavior inherent to the physical system.

As I mentioned in the Introduction, the driving field we have used throughout this thesis when examining absorptive bistability is actually above the bifurcation point for the related semi-classical system (the Maxwell-Bloch equations). The structure of this transition is still unknown in the quantum case, so I ran the LTSA algorithm on trajectories with driving fields ranging from $\mathcal{E} = 0.4$ to $\mathcal{E} = 0.65$. Over this range, the Q functions for the steady state of the quantum system transition from entirely single-peaked around the lower state, to entirely peaked around the upper state. The $d = 3$ LTSA manifold at $\mathcal{E} = 0.4$ consists of a cluster at the lower state, with its own shape and relation to system observables; as the driving field is increased, a second cluster appears and becomes more populated, eventually outgrowing the lower state near $\mathcal{E} = 11.2$, our canonical case. This process continues as the driving field \mathcal{E} increases, and at $\mathcal{E} = 0.65$, the upper state is left alone, with the LTSA manifold displaying the leading dimensions of its internal structure. Unfortunately, this process shows no fundamental or sharp change in the manifold topology corresponding to the bifurcation, and leaves us without significant additional insight into the bifurcation process.

4.4.1 Fitting to observables

As discussed above, the leading two dimensions from the LTSA algorithm applied to absorptive bistability are roughly monotonic in $\langle x \rangle$ and $\langle y \rangle$, which gives us hope that we will be able apply the same basis pursuit technique developed for phase bistability to these dimensions. The mystery of the third and fourth dimensions might be resolved by the search through variable space which basis pursuit automates, and the fifth (or sixth) dimension is likely related to σ_y .

Unfortunately, the possible basis set of first-, second- order polynomials of the first-, second-, and third-order combinations of our basic set of Hermitian system operators (x , y , σ_x , σ_y , and σ_z)

only do a good job of fitting the leading dimension when a large, unwieldy, number of the basis functions are included. The fit is generally better in the two stable regions, but becomes poor in the transition region as we try to restrict the number of number of basis states to fewer than 10 or so. Most likely, future work would require a larger basis extending beyond simple polynomials (and might require more computing power and/or more powerful fitting software), or external insight into the best functional form. While polynomials do span the space of possible functions, we may have lost considerably by considering only Hermitian polynomial terms independently.

The second dimension is well fit by

$$\hat{v}_2 = -0.0016 - 0.0178 \langle y \rangle - 0.004 \langle \sigma_z \rangle + 0.0014 \langle y \rangle \langle y^2 \rangle \quad (4.43)$$

for our example amplitude quadrature measurement data set with 2000 samples, $d = 5$, and $k = 40$. The leading contributors to the fit for the fifth dimension, as expected, relate to σ_y : $\langle \sigma_y \rangle$, $\langle \sigma_y \rangle \langle x \rangle$, $\langle \sigma_y \rangle \langle y^2 \rangle$, and $\langle \sigma_y \rangle \langle \sigma_z \rangle$. However, this set does only a marginal job at fitting this dimension alone, and several additional operators are necessary to get a high quality fit. Even with over a dozen terms contributing to each, the fit to the third and fourth dimensions are quite poor. Without a reasonable functional form for the leading dimension, we are unable to construct the equivalent of the new “Maxwell-Bloch” equations we constructed for phase bistability in Section 4.3.1. Regardless of this (perhaps temporary) obstacle, it is clear that Local Tangent Space Alignment provides valuable insight into the fundamental behavior of our example quantum systems.

Chapter 5

Conclusion

In this thesis, we have explored two possible techniques for developing simple models for the dynamics of quantum systems: Proper Orthogonal Decomposition (POD) and Local Tangent Space Alignment (LTSA). Through the lens of two particular regimes in cavity QED, we have evaluated these techniques and explored their advantages and disadvantages for quantum dynamical systems in general.

The main advantage of POD is that the subspace produced by the process is linear, which allows us to easily project the full master equation onto the space and develop filtering equations for the density matrix as a whole. However, this same linearity proves to limit the reduced models, as they are unable to capture the large deviations from the mean that characterize the nonlinear systems. Local Tangent Space Alignment produces a point-wise constructed nonlinear manifold, which can capture the nonlinear behavior. I was able to discover low-dimensional manifolds which capture the essential structure of the system dynamics. However, the lack of a functional form for the map between the manifold space and the density matrix space complicates the search for a reduced filter. In the case of phase bistability, I was able to fit the manifold coordinates by a linear combination of (linear and nonlinear) system observables, in order to derive a system-specific set of “Maxwell-Bloch” equations, which make a very good filter. This sort of fit, however, depends on having a good basis set of functions of observables with which to fit the data, which proved to be missing in the case of absorptive bistability. A structured way to create a much larger basis set might lead to success here, and in models of other quantum systems.

The model reduction techniques discussed here may also be applied in a wide range of systems, including other regimes from cavity QED as well as other nonequilibrium systems. The semi-classical Maxwell-Bloch equations display a range of behaviors in addition to the phase and absorptive bistabilities discussed in this thesis. As discussed by Armen and Mabuchi [22], they also have limit cycles related to Hopf bifurcations. POD and LTSA could be applied to the quantum trajectories which correspond to this semi-classical situation, and produce useful reduced models. The mean (and steady) state of the limit cycle takes the shape of a ring. The POD eigenstates are then

a progression of Fourier decompositions on the ring: First, a state whose Q function is a single positive zone and a single negative zone across from it. Then, a similar state whose features are rotated 90 degrees relative to the first state. Following this is a pair of states with two positive and two negative regions, rotated 45 degrees with respect to each other, then a pair with three positive and three negative, rotated 30 degrees, etc. Appropriate addition and subtraction of these sine- and cosine-like states can build a localized state, provided the basis is large enough. The projected filter on this basis, when driven by a white noise innovation, reproduces oscillation around the limit cycle. LTSA applied to the limit cycle produces a ring-shaped manifold, which requires at least two dimensions to contain. LTSA does not handle self-intersecting manifolds well, and it fundamentally cannot recreate a one-dimensional ring with $d = 1$.

The cavity QED system is only one of a large class of non-equilibrium systems where complex topology meets dynamics; other examples include chemical reaction dynamics and protein folding. Provided simulations (or experimental data) with sufficient resolution and coverage of the full space explored by these dynamics, POD or LTSA could draw out features of the underlying dynamical space. Nonlinear manifold learning algorithms such as LTSA are particularly exciting in this context, as these problems are generally interesting precisely because of their nonlinear features. (Understanding protein folding would be a very different problem if linear analysis were sufficient.) If we could construct a global functional form for the maps between the LTSA-generated submanifold and the large original data-space, projecting the dynamical equations to build simple filtering equations might provide significant insight into the dominant and critical features of these dynamics.

Future work to extend POD or LTSA (or adapt other algorithms) to build better filters for quantum systems might benefit by taking advantage of additional information present in the trajectory. Both POD and LTSA use the quantum trajectory as an unordered data set (although the zoned POD algorithm developed here is, in effect, time dependent). Fundamentally, however, we might expect that an algorithm which used the time ordering of the data would have an advantage. For example, a manifold learning algorithm might ensure that points which were nearby in time are also nearby on the manifold (provided the sampling is fine enough). This might also require a different definition of distance, which the failures of POD imply might be necessary in any event.

We have evaluated the performance of these model reduction techniques in part based on their ability to provide a simple filtering equation, which could be used in a control context. However, a controlled system might behave quite differently from the autonomous systems we have observed to build reduced models. In particular, I have not studied the input characteristics of the cavity QED or other quantum models, only their outputs. There is an extensive literature on balanced truncation [43] and other balancing transformations, such as [31], which make use of the input-output characteristics of the system, rather than (or in addition to) its internal structure and dynamics, to build reduced models.

Filtering equations are useful outside of a control context as well, because they allow us to simulate the dynamics of a system by driving it with white noise as the innovation. An exact filter, such as the full master equation, allows us to simulate the full quantum system; approximate filters give us insight into the system by giving us a simpler model which shows the same behaviors. If the approximate model has captured the essence of the exact system, we should be able to drive it with a white noise innovation process and see the approximate system reproducing the behavior of the exact one. We can then analyze the simple model for insight into the more complicated exact system. For example, the success of the “Maxwell-Bloch” equations derived in Section 4.3.1 indicates that these three (or five) equations contain the same kinds of interactions between nonlinearity and noise which also determine the dynamics of a quantum system in the phase bistable regime.

Filters produced by projecting the exact filter onto a lower-dimensional manifold (linear or nonlinear) may bring with them multiplicative noise terms, particularly if those terms are also present in the exact dynamics. Dynamical systems with multiplicative noise have two different kinds of special points, if they exist: points where the deterministic dynamics are stable, and points where the noise terms vanish (or reach a minimum in amplitude). Recall that van Handel and Mabuchi [11] showed that the phase bistable system reduces to the Wonham filter [44] in one limit. The stochastic terms of the Wonham filter vanish at the two points between which the system is observed to switch, while the deterministic terms have a stable point at the middle, where the system spends almost no time. Quantum bistable systems with multiplicative noise, and their high quality filters, may show similar behavior. By constructing multiplicative noise terms in approximate filters directly from the exact dynamics, the projected filter can correctly reproduce switching behavior which *ad hoc* or additive noise models cannot.

A complete understanding of the dynamics of open quantum systems will require understanding a variety of regimes and behaviors. The utility of open quantum systems depends in large part on the dynamics, such as switching, limit cycles, or bifurcation, which occur in nonlinear models. We have chosen to focus on these features of nonlinear systems, rather than chaotic behavior or other hallmarks of nonlinearity, because of their connection to the engineering of useful devices. Proper Orthogonal Decomposition and Local Tangent Space Alignment are just two from a large class of model reduction strategies applicable to these systems, and this thesis has evaluated them only in two parameter regimes of a single example system. However, we have been able to produce simple models which accurately capture interesting dynamics, while building a scaffold on which future work may build.

Bibliography

- [1] A. Hopkins, K. Jacobs, S. Habib, and K. Schwab. Feedback cooling of a nanomechanical resonator. *Phys. Rev. B*, 68:235328, 2003.
- [2] A. Hopkins, B. Lev, and H. Mabuchi. Proposed magnetoelectrostatic ring trap for neutral atoms. *Phys. Rev. A*, 70:053616, 2004.
- [3] C. J. Hood, M. S. Chapman, T. W. Lynn, and H. J. Kimble. Real-time cavity QED with single atoms. *Phys. Rev. Lett.*, 80(19):4157–4160, 1998.
- [4] J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, and H. J. Kimble. Deterministic generation of single photons from one atom trapped in a cavity. *Science*, 2004.
- [5] R. Miller, T. E. Northup, K. M. Birnbaum, A. Boca, A. D. Boozer, and H. J. Kimble. Trapped atoms in cavity QED: coupling quantized light and matter. *J. Phys. B: At. Mol. Opt. Phys.*, 38(9):S551–S565, 2005.
- [6] J. M. Raimond, T. Meunier, P. Bertet, S. Gleyzes, P. Maioli, A. Auffeves, G. Nogues, M. Brune, and S. Haroche. Probing a quantum field in a photon box. *J. Phys. B: At. Mol. Opt. Phys.*, 38(9):S535–S550, 2005.
- [7] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. Single-atom single-photon quantum interface. *Science*, 317(5837):488–490, 2007.
- [8] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963.
- [9] H. Mabuchi and H. M. Wiseman. Retroactive quantum jumps in a strongly coupled atom-field system. *Phys. Rev. Lett.*, 81:4620, 1998.
- [10] S. Tan. Quantum optics and computation toolbox for MATLAB (software), <http://www.qo.phy.auckland.ac.nz/qotoolbox.html>, October 2008.
- [11] R. van Handel and H. Mabuchi. Quantum projection filter for a highly nonlinear model in cavity QED. *J. Opt. B: Quantum Semiclass. Opt.*, 7:226, 2005.

- [12] L. Lugiato. Theory of optical bistability. In E. Wolf, editor, *Progress in Optics*, volume XXI. North-Holland, Amsterdam, 1984.
- [13] L. Perko. *Differential Equations and Dynamical Systems*. Springer-Verlag, New York, 2001.
- [14] L. Arnold. *Random Dynamical Systems*. Springer-Verlag, Berlin, 1998.
- [15] H. Kushner. *Stochastic Stability and Control*. Academic Press, New York - London, 1967.
- [16] R. Z. Khasminskii. *Stochastic Stability of Differential Equations*. Kluwer Academic, 2nd edition, 1980.
- [17] P. Alsing and H. J. Carmichael. Spontaneous dressed-state polarization of a coupled atom and cavity mode. *Quantum Optics*, 3:13, 1991.
- [18] S. Y. Kilin and T. B. Krinitskaya. Single-atom phase bistability in a fundamental model of quantum optics. *J. Opt. Soc. Am. B*, 8:2289–2295, 1991.
- [19] H. J. Carmichael. *An open systems approach to quantum optics : lectures presented at the Université libre de Bruxelles, October 28 to November 4, 1991*, volume m18 of *Lecture notes in physics*. Springer-Verlag, Berlin, 1993.
- [20] H. Gang, C. Ning, and H. Haken. Codimension-two bifurcations in single-mode optical bistable systems. *Phys. Rev. A*, 41(5):2702–2711, 1990.
- [21] H. Gang, C. Ning, and H. Haken. Distribution of subcritical hopf bifurcations and regular and chaotic attractors in optical bistable systems. *Phys. Rev. A*, 41(7):3975–3984, 1990.
- [22] M. Armen and H. Mabuchi. Low-lying bifurcations in cavity quantum electrodynamics. *Phys. Rev. A*, 73:063801, 2006.
- [23] A. Szoke, V. Daneu, J. Goldhar, and N. A. Kurnit. Bistable optical element and its applications. *Appl. Phys. Lett.*, 15(11):376–379, 1969.
- [24] C. M. Savage and H. J. Carmichael. Single atom optical bistability. *IEEE J. Quant. Elec.*, 24(8):1495–1498, 1988.
- [25] C. W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*. Springer, 2nd edition, 1985.
- [26] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. North Holland, 3rd edition, 2007.
- [27] P. E. Kloeden, E. Platen, and H. Schurz. *Numerical Solution of SDE Through Computer Experiments*. Springer, 1997.

- [28] J. L. Lumley. The structure of inhomogeneous turbulence. In A. M. Yaglom and V. I. Tatarski, editors, *Atmospherical turbulence and wave propagation*, pages 166–178. Nauka, Moscow, 1967.
- [29] K. Pearson. On lines and planes of closest fit to points in space. *Philosophical Magazine*, 2:609–629, 1901.
- [30] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24:417–441, 1933.
- [31] S. Lall, J. E. Marsden, and S. Glavaski. Empirical model reduction of controlled nonlinear systems. In *Proc. 1999 IFAC World Congress*, 1999.
- [32] Z. Zhang and H. Zha. Principal manifold and nonlinear dimensionality reduction via tangent space alignment. *SIAM J. Sci. Comput.*, 26(1):313–338, 2004.
- [33] A. Perelomov. *Generalized Coherent States and Their Applications*. Birkhäuser, 1986.
- [34] C. A. Fuchs and J. van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Trans. Info. Theory*, 45(4), 1216-1227 1999.
- [35] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319, 2000.
- [36] D. Freedman. Efficient simplicial reconstructions of manifolds from their samples. *IEEE Trans. Pattern Anal. and Mach. Intel.*, 24(10):1349, 2002.
- [37] D. L. Donoho and C. Grimes. Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. *Proc. Natl. Acad. Sci.*, 100(10):5591–5596, 2003.
- [38] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [39] L. K. Saul and S. T. Roweis. Think globally, fit locally: Unsupervised learning of low dimensional manifolds. *J. Mach. Learning Research*, 4:119–155, 2003.
- [40] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Appl. Comput. Harmon. Anal.*, 21:113–127, 2005.
- [41] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [42] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming (web page and software), <http://www.stanford.edu/~boyd/cvx/>, October 2008.

- [43] B. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Trans. Auto. Control*, 26(1):17–32, Feb 1981.
- [44] W. M. Wonham. Some applications of stochastic differential equations to optimal nonlinear filtering. *J. SIAM, Series A: Control*, 2(3):347–369, 1965.