Dynamics of an Ocean Energy Harvester

by

Clark McGehee

Department of Mechanical Engineering and Materials Science
Duke University

Date: _____

Approved:

_____
Brian P. Mann, Supervisor

_____
Laurens E. Howle

_____
W. Neal Simmons

_____
Richard D. Vann

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Mechanical Engineering and Materials Science
in the Graduate School of Duke University
2013

Abstract

Dynamics of an Ocean Energy Harvester

by

Clark McGehee

Department of Mechanical Engineering and Materials Science
Duke University

Date: _____

Approved:

_____
Brian P. Mann, Supervisor

_____
Laurens E. Howle

_____
W. Neal Simmons

_____
Richard D. Vann

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Mechanical Engineering and Materials Science
in the Graduate School of Duke University
2013

Copyright © 2013 by Clark McGehee
All rights reserved. The views and opinions expressed in this dissertation are solely those of the author and do not necessarily represent or reflect the official policy or position of the United States Air Force, Department of Defense or any agency of the United States government.

Abstract

Ocean-based wireless sensor networks serve a variety of important purposes, including monitoring and detection of tsunamis and seismic events, tracking the migration of endangered species, drug interdiction, and anti-submarine warfare. Developing energy harvesting devices that make these networks self-sufficient allows for reduced maintenance cost and greater reliability. Many methods exist for powering these devices, including internal batteries, photovoltaic cells and thermoelectric generators, but the most reliable method, if realized, would be to power these devices with an internal kinetic energy harvester capable of reliably converting wave motion into electrical power. Designing such a device is a challenge, as the ocean excitation environment is characterized by shifting frequencies across a relatively wide bandwidth. As such, conventional, linear kinetic energy harvesting designs are not capable of reliably generating power. Instead, a nonlinear device is better suited to the task, and the objective of this dissertation is to investigate the behaviors of devices that could be employed to this end.

The purpose of this dissertation is to build a comprehensive set of analysis tools that may be used in the design of an ocean energy harvester, and to apply those tools in the preliminary design optimization of a nonlinear energy harvester design. The first component of this dissertation focuses on the development of theoretical and computational techniques that can be used to accurately and efficiently characterize the response of a strongly nonlinear device through guided exploration of

the parameter space. The second half of the dissertation applies these techniques in conjunction with statistical methods to conduct a stochastic design optimization of a strongly nonlinear energy harvesting device for various sea states.

The first half of this dissertation focuses heavily on the theory of approximation methods and the development of computational tools to efficiently analyze the results of these methods. These techniques are applied in the analysis of the torque-excited pendulum and the nonlinear equations of motion of a piezoelectric inertial generator as illustrative examples. Using the example of a torque-excited pendulum, methods for obtaining approximate representations of the dynamical behavior of a rotating system are investigated and new insights regarding the inefficacy of conventional simplifying assumptions are found. It is shown that the conventional simplifying assumption of harmonic behavior is inadequate in the study of a rotating system, and an alternative method is presented to capture the rotating dynamics. Using the nonlinear energy harvester equations as an example, numerical tools are developed to conduct guided explorations of the parameter space and numerically quantify the parametric uncertainty of the results, eliminating the need for complex symbolic expressions of uncertainty and the reliance on closed-form solutions for the analysis of these systems.

In the second half of the dissertation, these techniques are applied in the study of a proposed nonlinear energy harvesting device. A phenomenological investigation and parametric study of the proposed device are conducted, and the insights drawn from the behavior of the system in simplified conditions are used to conduct a quasi-Monte Carlo investigation of the device in various sea states and to conduct a stochastic optimization of the design parameters to provide a first-order working estimate for the construction of a physical experiment.

In loving memory of my father, and of his jokes about the Thermos keeping hot stuff hot and cold stuff cold.

Contents

List of Tables

List of Figures

Acknowledgements

I owe an enormous debt of gratitude to many people who have helped me along my way as I have pursed my PhD. I have had the opportunity to work alongside people of tremendous caliber, and I will remember this experience fondly.

As I began the transition to college life, Dr. Michael R. Gustafson helped spur my interest in engineering by teaching a wonderful freshman engineering experience class at Duke University. Year after year, he continues to pour massive amounts of time and energy into the course to ensure that all freshman engineers have the best possible experience as they transition into the undergraduate engineering experience. I and many others are in his debt for the solid foundation he provided for our engineering careers.

My journey towards the doctoral program began with Dr. Brian Mann and Dr. Neal Simmons. Dr. Simmons' teaching style and love for the discipline ignited my interest in mechanical engineering and spurred my interest in engaging in mechanical engineering research. As part of his Engineering Innovation course, I was introduced to Dr. Mann, who took enormous amounts of time out of his schedule to explain his research and introduce me to his lab. Dr. Mann then assumed the responsibility for three Pratt Fellows students to ensure that I, as well as my classmates Ben Owens and Genevieve Lipp, were able to participate in his laboratory as undergraduate researchers. This went above and beyond what is expected, as few professors even take one Pratt Fellows student. Dr. Mann patiently instructed all of us and taught

us the skills we needed to succeed as researchers. It speaks to his character and ability that all three of us have since entered graduate school under his advising despite any other offers that have come our way.

Over the course of graduate school, Dr. Mann has been a constant source of motivation and expertise, but he has also been an excellent friend. I probably owe an apology to Mrs. Mann for introducing him to triathlons, but over the past several years he has become an avid triathlete and an excellent race partner. Somehow, he still manages to juggle a full teaching load, research work and a family life while still excelling at the sport. I look to his example as someone who knows how to get things done while staying in touch with the real world.

Dr. Laurens Howle has also been a constant collaborator since I entered graduate school. After enduring my nonsense as my undergraduate advisor, he for some reason decided that collaborating on financial markets research was a good idea. I am thankful for whatever led to this decision. I have learned a tremendous amount about a field that I had no knowledge of before graduate school under his advising, and I have learned much about the art of computer programming from his expertise. Without the skills I learned from our collaboration, I surely would not have been capable of handling the technical numerical work in this dissertation. He is yet another example of someone who knows how to get things done, but remains grounded in the practical aspects of the discipline, and I look to his example as I pursue my P.E. license and my future as an engineer.

Dr. Samuel Stanton was both an excellent collaborator and friend as I began my graduate studies. I appreciate the time and effort he spent teaching me about his research and the opportunities he gave me to participate in two major energy harvesting investigations.

I have had many wonderful Air Force colleagues who have also helped me along my way. Lt Col (ret.) JD Wroth served as the commander of Detachment 585

when I was a cadet in AFROTC. He always recognized people as a commander's most valuable asset, and he frequently went above and beyond to help his cadets. Without his help, I would not have been able to go to Duke. He set an example for leadership that I will strive to follow as I begin my career as an Air Force officer. I am thankful that Lt Col (ret.) Peter Oertel arranged my educational delay status so that I could pursue this degree. Dr. Jon Protz has offered excellent career advice and spurred me to find a way participate in the PhD program when I thought that all doors were closed. I am very thankful for his encouragement and for all of the advice he has given me.

Col (ret.) Harry Foster was a formational influence on my career as I transitioned from undergraduate to graduate school. Working with him his colleagues at the Center for Strategy and Technology was a unique experience that taught me much about the Air Force in a very short time. His continued advice and mentoring has been extremely helpful as I transition into the active duty Air Force.

Dr. Ronald Joslin's support as a grant manager from the Office of Naval research has made all of this possible. I appreciate his sponsorship and financial assistance for the past three years.

I would certainly not have come this far without the help of a loving family. I am very thankful for my parents, Rob and Cindy McGehee, who have given so much of their time and energy in helping me be where I am today. I am blessed to have had a mother who spent hours teaching me outside of the classroom and a father who worked incredibly hard to provide the best possible life for his family. It is doubtful that they will ever truly appreciate the impact they have had on my life. And finally, I am very thankful for Elizabeth, for the support she has provided, and for the future we will have together.

1

Introduction

Ocean-based autonomous wireless sensor networks serve a variety of important functions in scientific research, civil disaster prevention and national security. These sensor networks are typically comprised of tethered or untethered buoys, ranging in size from several meters in diameter in the case of tsunami warning buoys to less than a quarter of a meter in the case of sonobuoys. Such networks have been employed in tasks including the study of whale migration, observation of weather patterns, tracking and detection of submarines, and providing early warning for tsunami events. Each node in the sensor network collects and transmits data, expending electrical energy in the process. While advances in low-power sensors have led to devices that expend very little energy in collecting and transmitting data, the energy reserves of a sensor node remain finite in the absence of an energy capture mechanism. Developing such a mechanism is the focus of this work.

Figure 1.1 shows a schematic of a typical wireless sensor network with which this research is concerned. Such a network would typically be comprised of many small, untethered buoys with low-powered sensing equipment. Photovoltaic cells may be used to power some of the device components but in certain operating environments

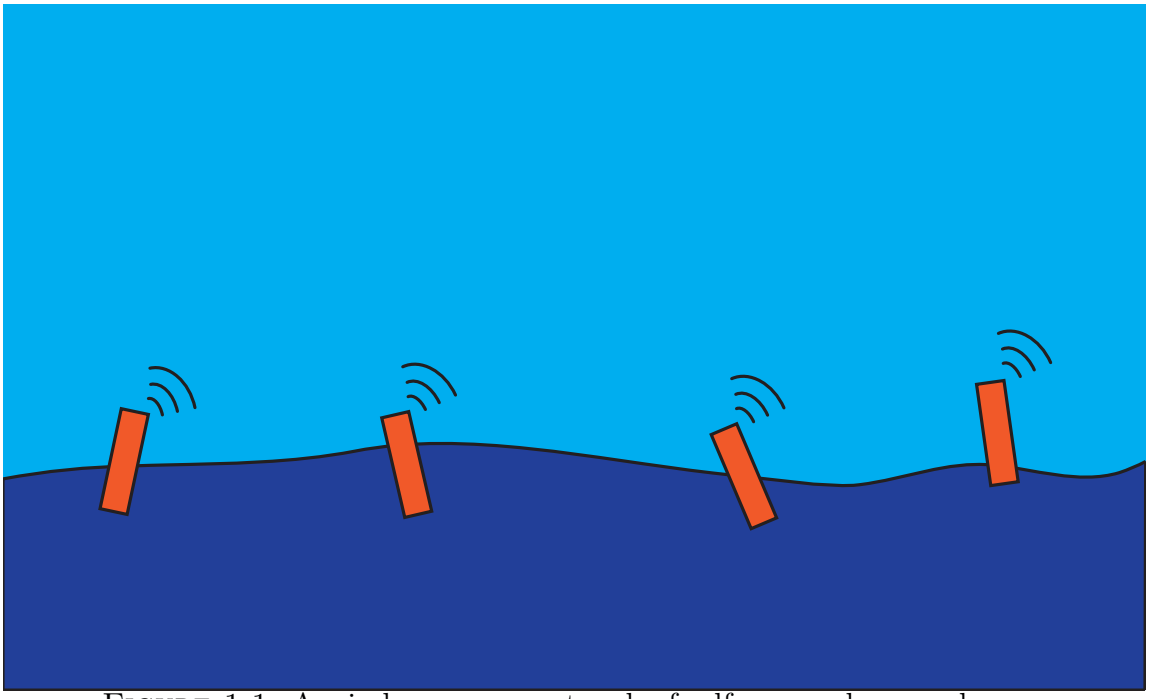

FIGURE 1.1: A wireless sensor network of self-powered sensor buoys.

solar power is too unreliable to provide a reliable source of power. An alternative to solar power is to extract electrical power from the kinetic energy of the ocean waves through an electromechanical device known as an energy harvester. Kinetic energy harvesters take many forms, ranging from magnet-coil designs [2] to piezolaminated beams [3], but fundamentally, their operating premise is the same: the device converts mechanical energy to electrical energy through a transduction mechanism such as electromagnetism or piezoelectricity. Fundamental to the understanding of these devices is the notion of coupling; any electromechanical system exhibits coupling behavior between its electrical and mechanical components. Coupling behavior can effect the dynamics of the system in unexpected ways [2, 3, 4, 5], so it is important to consider the dynamic effects of coupling in any energy harvesting investigation.

Harvesting energy from ocean waves presents additional challenges due to the excitation environment. Figure 1.2 illustrates the nature of the ocean spectral environment. The spectrum is comprised of a relatively broad band of low frequencies

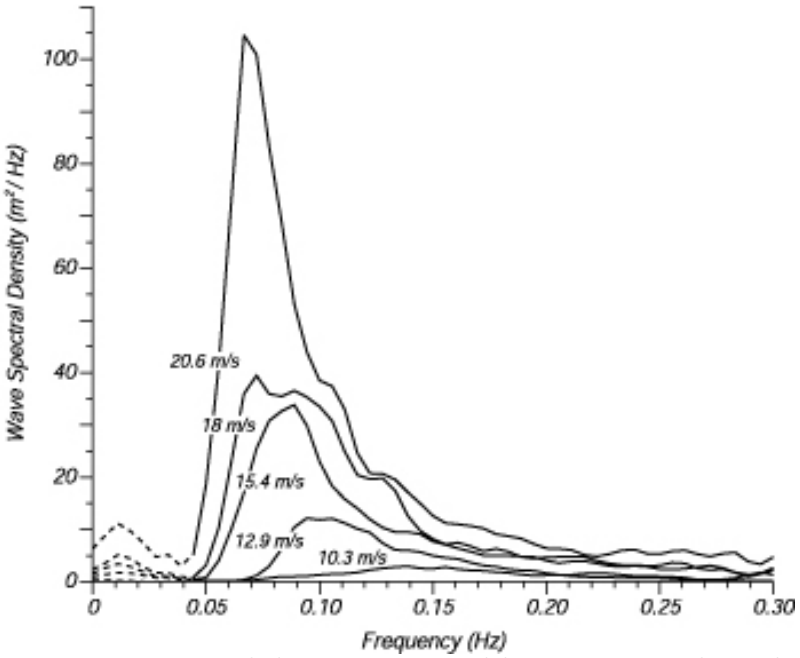

FIGURE 1.2: Wave spectral data as reported by Pierson and Moskowitz [1].

that vary with wind speeds. The shifting nature of the frequency content is the primary obstacle in designing a device to capture energy from this spectrum. It has been demonstrated in the literature that nonlinear devices are well-suited to the task of energy capture in broadband, shifting frequency environments [2, 3, 4]; thus a nonlinear device for kinetic energy capture is proposed and analyzed in this dissertation.

Figure 1.3 shows a conceptual drawing of an energy harvesting device intended for ocean applications. Inspired by the self-winding wristwatch, the device relies on its ability to collect the intermittent bursts of energy provided by ocean waves to maintain energy in the connected electrical energy storage devices. The stored energy is then used to operate the sensors and transmit collected data.

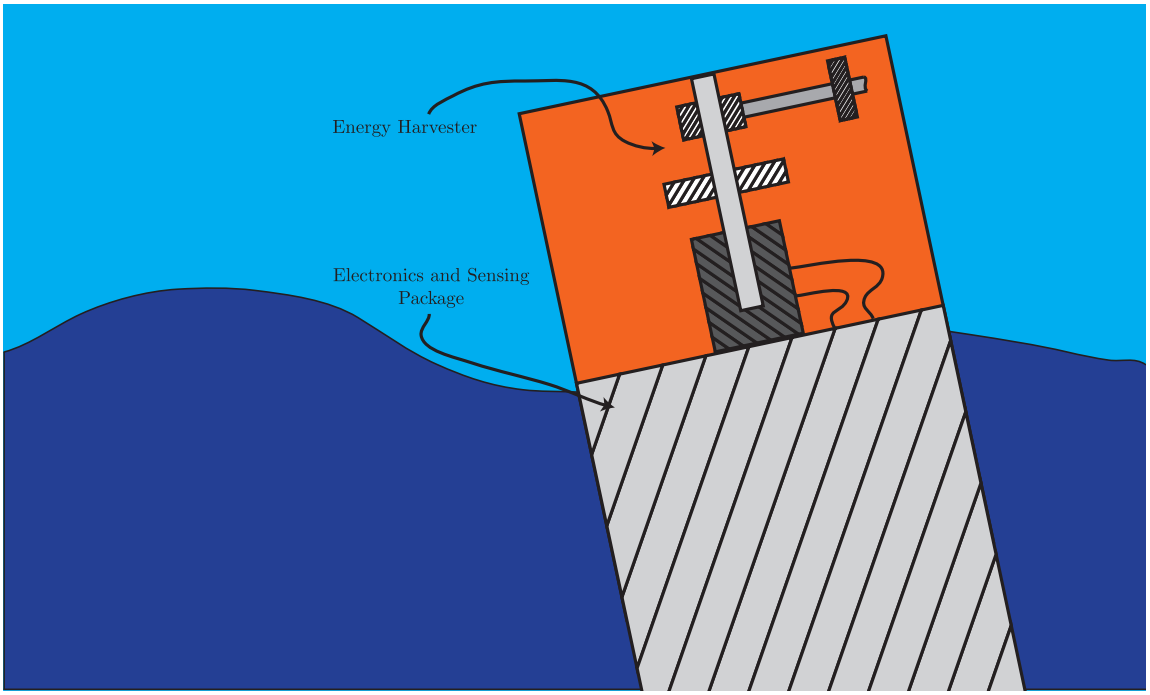

FIGURE 1.3: Conceptual design of an ocean energy harvester in a wireless sensor buoy.

1.1 Research contributions

The primary objective of this research is to advance the process of designing an ocean energy harvester from the conceptual design shown in figure 1.3 to a set of concrete design recommendations for a prototype of the device that could be employed in an ocean environment. In the course of studying this problem, several tangential discoveries were made, which are detailed within this section.

1.1.1 Detailed comparison of harmonic and anharmonic generating solutions applied to the torque-forced pendulum

Chapter 3 contributes an improved strategy for averaging strongly nonlinear systems and provides a detailed justification for choosing the new approach. The work presents a qualitative comparison of two generating solutions, shown in figure 1.4, as applied to the method of averaging using the example of a torque-excited pendulum.

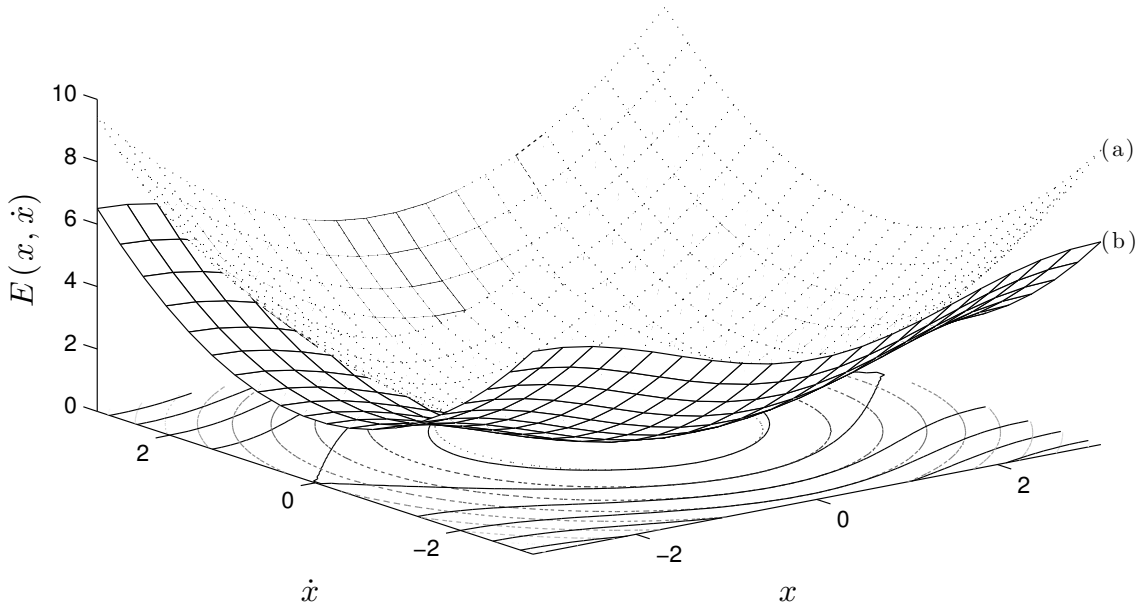

FIGURE 1.4: Contours of harmonic (a) and anharmonic (b) generating solutions and potential functions.

Prior work by Yuste [6] conducts a quantitative comparison of similar methods in the context of generalized harmonic balance; this work adds a visual element to the argument and clearly demonstrates the importance of selecting an appropriate generating solution when using the method of averaging. In the analysis of a torque-excited pendulum, one may opt for a harmonic generating solution or an anharmonic generating solution derived from the solution of the unforced pendulum equations. It is shown that the harmonic generating solution provides poor results in comparison to the anharmonic generating solution; in particular, the harmonic solution suffers from accuracy problems at large amplitudes and does not capture the rotating solutions that occur for higher amplitude forcing near resonance. In contrast, the anharmonic solution compensates for the amplitude dependence of the solution period, increasing accuracy at high amplitudes, and it captures the rotating solution organically.

1.1.2 Efficient numerical methods for the analysis and uncertainty quantification of nonlinear energy harvester models

Chapter 4 contributes a simplified method for parameter studies and uncertainty quantification of the nonlinear energy harvester models of references [7, 3, 8, 4]. To date, the energy harvesting community has not embraced the use of numerical methods beyond numerical integration of the ordinary differential equation models in the energy harvesting literature, and most parameter studies are conducted using analytical (symbolic) methods to find a closed-form approximation of the periodic orbits of the system. A major drawback to these approaches are extreme computational expense for detailed studies of the parameter space and unnecessarily complex symbolic computations, respectively. Computing the periodic orbits by psuedospectral collocation and employing the techniques of numerical continuation and bifurcation analysis surmounts both of these challenges by greatly reducing computational effort and removing the need for unwieldy analytical solutions. Additionally, conducting uncertainty quantification numerically avoids additional complex symbolic computations and can be performed at minimal additional computational cost.

1.1.3 Derivation and study of a mathematical model for a horizontal pendulum

Chapter 5 presents a novel device with interesting physical behaviors with a proposed application as the mechanism of kinetic energy capture in an ocean-based energy harvester. The device is inspired by the self-winding wristwatch, a proven energy harvesting technology that has been in use for nearly a century [9]. The chapter derives an equation of motion for the system and conducts a phenomenological and parametric study of its behaviors. Insights are drawn and applied to the design analysis in the subsequent chapter.

1.1.4 Design analysis and optimization of a proposed energy harvester design

Chapter 6 utilizes a geostatistical approach to conduct a study of the parameter space for the device proposed in chapter 5 in an ocean excitation environment. A representation of deep-water ocean random waves is presented, and analytical expressions for the surface angle of the water and its temporal derivatives are presented and used to numerically simulate the behavior of the system. A quasi-Monte Carlo investigation is conducted in the parameter space to construct a response surface, which is then fit by kriging (Gaussian process regression) [10]. This information is then used to make recommendations for the design of an energy harvesting prototype from this device.

1.1.5 Classification of price dynamics in equities markets

Chapter 7 focuses on an entirely different problem - characterizing the short-term behavior of electronically traded securities. In this investigation, a new reduced-order vector autoregression model is proposed and compared to an established model from the literature. It is demonstrated that price formation is largely driven by the forces of supply and demand on short time scales. The primary objective of this investigation was to find a reduced feature space that could be used to form a discriminative classifier for price movements; a feature space consisting of two demand variables was constructed and it was shown that price movements are clearly separated in the proposed feature space.

1.2 Organization of this dissertation

This dissertation is organized as follows: In chapter 2, a modeling framework for physical systems encountered within the dissertation is developed and discussed. Important theorems are proved and references are provided for the reader. Chapter 3 then employs the methods discussed in chapter 2 to investigate the dynamics of

the driven vertical pendulum. Both the librating and rotating solutions are studied, and the influence of the generating solution on the quality of results obtained is highlighted. Chapter 4 studies the numerical energy harvester equations using numerical methods. Parameter studies and uncertainty quantification are conducted numerically and it is demonstrated that numerical methods are capable of obtaining the same results as analytical methods without the need for a complicated analytical solution. Chapter 5 focuses on the development of a mathematical model for an ocean energy harvester, employing techniques from the prior chapters in the analysis of the device. Chapter 6 conducts a quasi-Monte Carlo investigation of the device under random forcing in a wave environment simulated by the Pierson-Moskowitz spectrum. Finally, Chapter 7 presents a study of the electronically traded fund QQQ in the NASDAQ market. Price formation on short time scales is studied and the law of supply and demand is observed even in the highly chaotic environment of the market microstrucutre.

2

Methods

Before presenting any investigations conducted in this dissertation, it is useful to develop a common modeling framework to give the reader a top-level view of the thought process behind each chapter. Despite any disparities in topical matters between the chapters, there is a consistent approach to the treatment of dynamical systems in both modeling and analysis throughout the dissertation. Beginning with a physical system, equations of motion are derived from first principles. These equations are investigated with numerical simulation, which allows for the visualization of dynamic behavior for a given set of parameters and initial conditions. Reasonable parameter values are typically determined by constructing an experiment and measuring or fitting the parameters of the experimental device to the model. To obtain a more global understanding of the system's behavior, it is then analyzed using either analytical or numerical methods such as the method of averaging or numerical continuation.

2.1 Mathematical models of physical systems

The first task in analyzing any complex system is to build a working model of the system. When modeling a system, several considerations must be made regarding the detail and complexity of the model. If too complex a model is chosen, analysis becomes prohibitively difficult, while if too simplistic a model is chosen analysis may provide misleading, oversimplified results. In practice we seek the optimal tradeoff between the two.

Mathematical modeling is an enormously general term that encompasses a wide swath of the applied mathematical community. The author recommends the M.H. Holmes texts [11, 12] for introductions to the foundations of applied mathematics and perturbation methods, respectively. The text by Bender and Orszag [13] is also highly recommended. Hinch [14] is another excellent reference on perturbation methods. For the statistical work within the dissertation, and in particular the work found in chapter 7, the machine learning text by Bishop [10] is the definitive reference.

The literature on mathematical modeling is quite broad, and it is necessary to narrow the focus of discussions on mathematical modeling to dynamical systems theory for the purposes of this dissertation. Nestled at the intersection of physics and mathematical modeling, the field of dynamical systems arose from the work of Henri Poincaré and his contemporaries in the late 19th and early 20th centuries as they sought to develop sophisticated mathematical tools to further the study of celestial mechanics. As the field has evolved in the past century, the tools and methods developed by Poincaré and others have been adapted for use in fields ranging from biology to quantum mechanics. Still, the most rigorous implementation of these theories is found in the physical sciences, providing scientists with the ability to model physical phenomena with a high degree of accuracy. Recommended texts include Greenwood's text on advanced dynamics [15] and Gelfand and Fomin's text

on the calculus of variations [16].

2.1.1 Hamilton's Principle or the Principle of Stationary Action

The standard approach to modeling physical systems taught in undergraduate courses is to balance external forces with inertia in accordance with Newton's Second Law[1],

$$\sum F_x = m\ddot{x}, \tag{2.1}$$

where the overdot ($\dot{}$) indicates a derivative with respect to time. By substituting mathematical representations of external forces derived from empirical observations and constitutive relationships, one arrives at an equation of motion in x that describes the motion of physical system in question.

While Newtonian mechanics provides an excellent framework for modeling simple physical systems, it quickly becomes unwieldy as the complexity of the physical system increases. Modeling of more complex systems is typically done using a framework known as Hamilton's principle, or the principal of stationary action.

Definition 2.1 (Principle of Stationary Action)**.** *The principle of stationary action states that the action* $\mathbb{S}[\boldsymbol{q}(t)]$*, defined as*

$$\mathbb{S}[\boldsymbol{q}(t)] = \int_{t_0}^{t_f} \mathcal{L}(\boldsymbol{q}(t), \dot{\boldsymbol{q}}(t), t)\,\mathrm{d}t \tag{2.2}$$

is stationary at $\boldsymbol{q}^*(t)$*, the true trajectory of the system. That is,* $\boldsymbol{q}^*$ *is the solution of the functional equation*

$$\frac{\delta \mathbb{S}}{\delta \boldsymbol{q}(t)} = 0. \tag{2.3}$$

[1] Note: this is the simplest possible formulation of N2; there exist others.

The interpretation of Hamilton's principle is that natural systems seek to expend the least energy possible in achieving any sort of reconfiguration: in other words, nature seeks the path of least resistance. This provides a rigorous mathematical framework under which the equations of motion of physical systems can be derived from first principles. Given expressions for the kinetic and potential energy of a system, one may use Hamilton's principle to determine equations of motion in an arbitrary set of coordinates.

A result of Hamilton's principle are the Euler-Lagrange equations, which form a necessary condition for the optimality of a trajectory as defined under the framework of stationary action. As an example of how the kinetic and potential energy of the system can be used to derive equations of motion, it is illustrative to derive the Euler-Lagrange equations, and provide a simple example.

Lemma 2.2 (Euler-Lagrange Equations)**.** *The Euler-Lagrange equations*

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathcal{L}}{\partial \dot{q}}\right) - \frac{\partial \mathcal{L}}{\partial q} = 0 \tag{2.4}$$

are a necessary condition for stationary action.

Proof. Denote the optimal action by

$$\mathbb{S}_* = \mathbb{S}[q_*] = \int_{t_0}^{t_f} \mathcal{L}(q_*, \dot{q}_*, t)\,\mathrm{d}t; \tag{2.5}$$

thus, the Lagrangian $\mathcal{L}_* = \mathcal{L}(q_*, \dot{q}_*, t)$ minimizes the action because $q_*(t)$ is the path of least action. Then a small perturbation from the optimal path, denoted as $\epsilon h(t)$, yields a Lagrangian $\mathcal{L}_\epsilon = \mathcal{L}(q_* + \epsilon h, \dot{q}_* + \epsilon \dot{h}, t)$ which results in an action greater than that of the optimal path, that is

$$\mathbb{S}[q_*] \geqslant \mathbb{S}[q_* + \epsilon h]. \tag{2.6}$$

A Taylor expansion of $\mathbb{S}[q_* + \epsilon h]$ yields

$$\mathbb{S}[q_* + \epsilon h] = \int_{t_0}^{t_f} \left\{ \mathcal{L}_* + \epsilon \left(\frac{\partial \mathcal{L}_*}{\partial q} h + \frac{\partial \mathcal{L}_*}{\partial \dot{q}} \dot{h} \right) + \mathcal{O}(\epsilon^2) \right\} \mathrm{d}t. \tag{2.7}$$

This equation can be broken into two parts by neglecting the higher order terms ($\mathcal{O}(\epsilon^2)$):

$$\mathbb{S}[q_* + \epsilon h] = \mathbb{S}[q_*] + \epsilon \int_{t_0}^{t_f} \left(\frac{\partial \mathcal{L}_*}{\partial q} h + \frac{\partial \mathcal{L}_*}{\partial \dot{q}} \dot{h} \right) \mathrm{d}t. \tag{2.8}$$

Note that the condition for optimality is

$$\frac{\delta \mathbb{S}}{\delta q} = \int_{t_0}^{t_f} \left(\frac{\partial \mathcal{L}_*}{\partial q} h + \frac{\partial \mathcal{L}_*}{\partial \dot{q}} \dot{h} \right) \mathrm{d}t = 0. \tag{2.9}$$

Integrating by parts yields

$$\frac{\delta \mathbb{S}}{\delta q} = \int_{t_0}^{t_f} \frac{\partial \mathcal{L}}{\partial q} h \, \mathrm{d}t + \left. \frac{\partial \mathcal{L}}{\partial \dot{q}} h \right|_{t_0}^{t_f} - \int_{t_0}^{t_f} \frac{\mathrm{d}}{\mathrm{d}t} \left(\frac{\partial \mathcal{L}}{\partial \dot{q}} \right) h \, \mathrm{d}t = 0; \tag{2.10}$$

utilizing the fact that the perturbation h vanishes at the endpoints yields

$$\frac{\delta \mathbb{S}}{\delta q} = \int_{t_0}^{t_f} \left\{ \frac{\partial \mathcal{L}}{\partial q} - \frac{\mathrm{d}}{\mathrm{d}t} \left(\frac{\partial \mathcal{L}}{\partial \dot{q}} \right) \right\} h \, \mathrm{d}t + \cancelto{0}{\left. \frac{\partial \mathcal{L}}{\partial \dot{q}} h \right|_{t_0}^{t_f}} = 0; \tag{2.11}$$

and given the arbitrary nature of h within the endpoints, the condition

$$\frac{\partial \mathcal{L}}{\partial q} - \frac{\mathrm{d}}{\mathrm{d}t} \left(\frac{\partial \mathcal{L}}{\partial \dot{q}} \right) = 0 \tag{2.12}$$

is necessary to satisfy $\delta_q \mathbb{S} = 0$. □

2.1.2 Example: Spring-Mass System

A simple example using the Euler-Lagrange equations to derive the equation of motion for a mass-spring system would proceed as follows. The kinetic and potential energies of the system are

$$\mathcal{T} = \frac{1}{2} m \dot{x}^2 \tag{2.13a}$$

$$\mathcal{U} = \frac{1}{2} k x^2 \tag{2.13b}$$

where m is the mass, and k is the spring constant. The variable x is known as a generalized coordinate, but in this case is simply the displacement of the mass attached to the spring relative to a fixed reference frame. The overdot $\dot{(\,)}$ is simply a derivative with respect to time. The Lagrangian is formed by subtracting the potential energy from the kinetic energy,

$$\mathcal{L}(x, \dot{x}, t) = \mathcal{T} - \mathcal{U} = \frac{1}{2} m \dot{x}^2 - \frac{1}{2} k x^2, \tag{2.14}$$

and can be substituted into the Euler-Lagrange equations to determine the equation of motion. Term by term, this yields

$$\frac{\partial \mathcal{L}}{\partial x} = -kx \qquad \text{restoring force terms} \tag{2.15a}$$

$$-\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathcal{L}}{\partial \dot{x}}\right) = -m\ddot{x} \qquad \text{inertial force terms} \tag{2.15b}$$

Substituting into equation (2.12) and multiplying both sides by -1 yields the familiar equation for an undamped, unforced mass-spring oscillator,

$$m\ddot{x} + kx = 0, \tag{2.16}$$

and illustrates why the Euler-Lagrange equations are typically written as in equation (2.4) and not as in equation (2.12). In this case, "Euler-Lagrange *equations*" is somewhat of a misnomer, as there is only one equation, but in the case of systems with more than one generalized coordinate there will be one equation for each coordinate, hence the naming convention. Note that this formulation only admits conservative terms arising from the exchange of kinetic and potential energy; more complex formulations of the Euler-Lagrange equations incorporate generalized forces and other nonconservative terms, but are not included here for the sake of brevity. The reader is directed to Greenwood [15] for an exhaustive discussion of variational modeling techniques.

2.2 Experimental parameter identification

Identifying a reasonable parameter set is an important step in any investigation. Typically, most parameters can be identified through measurement or are provided by the manufacturer's specifications for any off-the-shelf components. However, there are many cases when parameters must be determined experimentally. To do so, an experimental device is constructed and operated within the limits of an available analytical solution. Then, the analytical solution, which incorporates all of the model parameters as variables, is fit to the experimental data. By fitting the analytical solution to the experimental data, a best-fit set of parameters can be found that describe the experimental system.

2.2.1 Least-squares parameter estimation

Consider an experimental data set $\boldsymbol{y}(t)$. Typically, several (n) channels of experimental data would be sampled stroboscopically at a sampling frequency F_s for a

total of m samples. Thus the data set $\boldsymbol{y}$ could be represented as a matrix

$$\boldsymbol{y} = \begin{pmatrix} y_1(t_1) & y_2(t_1) & \cdots & y_n(t_1) \\ y_1(t_2) & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ y_1(t_m) & y_2(t_m) & \cdots & y_n(t_m) \end{pmatrix} \tag{2.17}$$

where each y_i for $i \in [1, n]$ is a channel and each t_j for $j \in [1, m]$ is a sample time. Now consider a candidate solution $\boldsymbol{x}(t) = \boldsymbol{\Phi}(t; \boldsymbol{p})$ that is hypothesized to describe $\boldsymbol{y}$; $\boldsymbol{p}$ is a vector of parameters of length $\leqslant m$ that describe the model. In the simplest case, $\boldsymbol{\Phi}$ is linear in $\boldsymbol{p}$. The classic example would be a polynomial series $b + p_1 t + p_2 t^2 \ldots$,

$$\boldsymbol{\Phi}(t)\boldsymbol{p} = \begin{pmatrix} t_1 & t_1^2 & \cdots & t_1^n & 1 \\ t_2 & \ddots & \ddots & \vdots & 1 \\ \vdots & \ddots & \ddots & \vdots & 1 \\ t_n & t_n^2 & \cdots & t_m^n & 1 \end{pmatrix} \begin{pmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \\ b \end{pmatrix}; \tag{2.18}$$

while the polynomials are certainly nonlinear, the coefficients are linear and thus techniques from linear algebra may be employed to solve the problem. This yields the optimization problem

$$\boldsymbol{p} = \arg\min_{\boldsymbol{p}} \sum_{j=1}^{m} \left(\boldsymbol{y}(t_j) - \boldsymbol{x}(t_j)\right)^2 = \arg\min_{\boldsymbol{p}} \left(\boldsymbol{y}^\top \boldsymbol{y} - 2(\boldsymbol{\Phi}\boldsymbol{p})^\top \boldsymbol{y} + (\boldsymbol{\Phi}\boldsymbol{p})^\top \boldsymbol{\Phi}\boldsymbol{p}\right). \tag{2.19}$$

By differentiating the residual with respect to $\boldsymbol{p}$ and setting it to zero,

$$\frac{\partial}{\partial \boldsymbol{p}} \left(\boldsymbol{y}^\top \boldsymbol{y} - 2(\boldsymbol{\Phi}\boldsymbol{p})^\top \boldsymbol{y} + (\boldsymbol{\Phi}\boldsymbol{p})^\top \boldsymbol{\Phi}\boldsymbol{p}\right) = \boldsymbol{\Phi}^\top \boldsymbol{\Phi}\boldsymbol{p} - \boldsymbol{\Phi}^\top \boldsymbol{y} = 0 \tag{2.20}$$

the optimal parameter set $\hat{\boldsymbol{p}}$ is found to be

$$\hat{\boldsymbol{p}} = \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^\top \boldsymbol{y} = \boldsymbol{\Phi}^+ \boldsymbol{y} \tag{2.21}$$

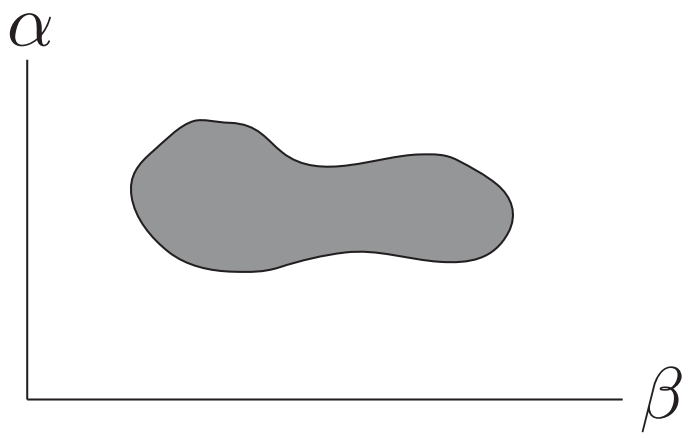

FIGURE 2.1: A sketch of a 2-D parameter space in the parameters α and β. The system exhibits a certain type of behavior within the shaded region that is desirable and thus in designing the system α and β are chosen to fall within the shaded region.

where $\left(\mathbf{\Phi}^\top \mathbf{\Phi}\right)^{-1} \mathbf{\Phi}^\top = \mathbf{\Phi}^+$, the Moore-Penrose psuedoinverse. This produces an optimal answer in ideal conditions (i.e. $m \gg$ the number of parameters); for methods to handle non-ideal conditions see [10].

In many cases, models used in vibrations will be nonlinear in the parameters; for example fitting the frequency ω of a harmonic function (e.g. $\sin \omega t$) is a nonlinear optimization problem. To fit these models, a nonlinear optimization routine such as the Levenberg-Marquardt method is recommended [17]. Those with access to MATLAB's Optimization Toolbox [18] should consider the routines `fminunc` and `fmincon` for unconstrained and constrained nonlinear optimization problems, respectively.

2.3 Analysis of mathematical models

A next step in the analysis of any device is to conduct parameter studies. By specifying a reasonable range for the device's parameters and obtaining an analytical or numerical representation of the behavior of the device, a parameter study can be conducted that allows the device to be optimized in the parameter space.

2.3.1 Analytical methods and the method of averaging

One such approach is to find an analytical, or symbolic, solution to the equations of motion. Such solutions are computationally advantageous: a closed-form solution eliminates the need for a numerical algorithm to find that solution at any point in the analysis process. However, they are often difficult to obtain and may require simplifying assumptions that degrade the quality of the analysis if the assumptions made are too broad.

A plethora of analysis methods for nonlinear ordinary differential equations exist, and the literature ranges from abstract to applied. As engineers, we tend to gravitate towards the applied literature. For an excellent introduction to the topic of nonlinear oscillations, the author recommends the text by Jordan and Smith [19] and accompanying sourcebook. Detailed worked examples are included, and many are directly relevant to nonlinear vibrations. A classic text in the field of nonlinear oscillations is the monograph by Guckenheimer and Holmes [20]. While less accessible as an introductory text, it is a thorough and detailed accounting of most of the mainstream topics in the field. The texts by Holmes [12] and Hinch [14] prove useful in understanding the perturbation method and asymptotic expansion techniques used in both works.

Method of averaging

One such analysis method, featured in chapter 3, is the method of averaging. The method of averaging is a mathematically rigorous way to replace a slowly-varying periodic vector field with its average, which often greatly simplifies the equations and transforms a previously intractable problem into a form that is amenable to analysis.

There are many excellent sources on the method of averaging, including the theory and mathematical fineries of the method. The author recommends the monograph by Sanders, Verhulst and Murdock [21] for a heavily theoretical background on the

subject from an applied mathematics perspective. The previously mentioned texts by Jordan and Smith [19] and Guckenheimer and Holmes [20] offer simpler worked examples. An additional text by Verhulst [22] offers a more casual discussion of the topic while retaining a high level of mathematical rigor.

Borrowing from [21], averaging is concerned with solving a perturbation problem of the form

$$\dot{\boldsymbol{x}} = \epsilon \boldsymbol{f}(\boldsymbol{x}, t) + \epsilon^2 \boldsymbol{g}(\boldsymbol{x}, t, \epsilon), \quad \boldsymbol{x}(0) = \boldsymbol{x}_0 \tag{2.22}$$

where $\boldsymbol{f}$ and $\boldsymbol{g}$ are T-periodic in t. Averaging solves the truncated ($\epsilon^2 \to 0$) system

$$\langle \dot{\boldsymbol{z}} \rangle = \epsilon \bar{\boldsymbol{f}}(\boldsymbol{z}), \quad \boldsymbol{z}(0) = \boldsymbol{x}_0 \tag{2.23}$$

where

$$\bar{\boldsymbol{f}}(\boldsymbol{z}) = \frac{1}{T} \int_0^T \boldsymbol{f}(\boldsymbol{z}, t)\, \mathrm{d}t. \tag{2.24}$$

As shown in [21], this yields an approximation $\boldsymbol{z}(t)$ to the solution of the full vector field $\boldsymbol{x}(t)$ with the properties

$$||\boldsymbol{x}(t) - \boldsymbol{z}(t)|| \leqslant C_1 \epsilon \quad \text{for} \quad 0 \leqslant t \leqslant C_2/\epsilon, \tag{2.25}$$

for positive constants C_1 and C_2. In other words, a solution that is asymptotically equal (error $\leqslant \epsilon$ times a constant C_1) to the true solution for a period of time $t \leqslant 1/\epsilon$ times a second constant C_2. The takeaway is that sources of error (i.e. the size of ϵ) should be minimized if one desires the most accurate result. Chapter 3 provides a more in-depth look at this topic.

2.3.2 Numerical continuation and the Implicit Function Theorem

Numerical continuation is an important tool in the analysis of dynamical systems. By finding an equilibrium point, the Implicit Function Theorem states that for small variations in parameter values, the position of the equilibrium should not change appreciably unless the system undergoes a bifurcation. This fact can be leveraged to create an algorithm that follows the equilibria of a dynamical system, the roots of a nonlinear equation, or even the periodic orbits of a nonlinear oscillator.

There is a wealth of excellent literature on continuation. The texts by Kuznetsov [23], Seydel [24] and Govaerts [25] are excellent starting points. The author was first introduced to and trained in numerical continuation through private correspondence with Dr. David Barton (University of Bristol) in collaboration with Dr. Mann. His expertise and assistance is greatly appreciated. Many software packages for continuation exist, including AUTO and MATCONT, however in this dissertation these packages are not used. For a detailed discussion of continuation methods used within this dissertation, refer to chapter 4.

2.4 Conclusions

This chapter presents a systematic approach to the modeling and analysis of physical systems encountered within this dissertation. It is hoped that the discussion and references provided within this chapter will aid the reader in understanding the modeling and analysis decisions made throughout the dissertation. There are of course a multitude of other methods that could be used to achieve the same ends, however, the methods listed herein are those preferred by the author in the modeling, simulation and analysis of nonlinear dynamical systems.

3

Comparison of Generating Solutions and the Resulting Averaged Vector Fields in the Study of the Torque-Excited Pendulum

In the course of investigating the horizontal pendulum system, it was discovered that the equation describing a damped vertical pendulum often arose during the process of transforming the system to a form amenable to averaging methods. Work stemming from analytical investigations of the horizontal pendulum system yielded new results regarding the efficacy of certain generating solutions in obtaining desirable results. In the study of strongly nonlinear systems, one must choose a generating solution for an averaging transformation that is capable of describing the behavior one seeks to model or that behavior will not be represented in the results. A brief quantitative argument to this effect regarding the generalized method of harmonic balance was presented by Yuste in reference [6], but to my knowledge there has been no such comparison for averaging methods presented in the literature. This chapter contributes a detailed qualitative study of harmonic and anharmonic generating solutions as applied to the analysis of the torque-excited pendulum, with the objective

of providing an intuitive visual message regarding the choice of generating solution to the reader through qualitative graphical comparisons.

3.1 Introduction

In the study of nonlinear vibrations, we often seek a closed-form approximation of both the transient and periodic steady-state behavior of a system using a variety of approximate analytical methods. One such method, averaging, functions by averaging the perturbations of a slowly moving vector field. The method of averaging, first proposed by Krylov and Bogoliubov[1] [26], is well studied and its mathematical properties are well understood [21]. However, the method itself cannot be directly applied to the equations most commonly encountered in nonlinear vibrations without first transforming from the state variables to amplitude-like and phase-like variables. Herein we show that the choice of transformation has a strong influence on the accuracy of the averaged solution. We consider the example of a torque-forced, viscously damped pendulum.

Conventionally, the method of averaging treats a nonlinear system as a slightly perturbed harmonic oscillator of the form

$$\ddot{x} + \omega^2 x = \epsilon f(x, \dot{x}) \tag{3.1}$$

where ϵ is a small parameter and ω is the linear natural frequency. This leads to to a natural transformation from the variables $(x, \dot{x}) \mapsto (r, \phi)$ as

$$\dot{r} = \epsilon f(r \sin \omega t, \omega r \cos \omega t) \omega r \cos \omega t \tag{3.2}$$

$$r\dot{\phi} = -\epsilon f(r \sin \omega, t, \omega r \cos \omega t) \sin \omega t \tag{3.3}$$

These equations are then averaged over one period of the unperturbed solution, i.e.

[1] Transliterations from the Cyrillic spelling may vary.

$$\dot{r} = \frac{\epsilon\omega}{2\pi}\int_0^{2\pi/\omega} f(r\sin\omega t, \omega r\cos\omega t)\omega r\cos\omega t \mathrm{d}t \tag{3.4}$$

$$r\dot{\phi} = -\frac{\epsilon\omega}{2\pi}\int_0^{2\pi/\omega} f(r\sin\omega, t, \omega r\cos\omega t)\sin\omega t \mathrm{d}t \tag{3.5}$$

and solved first for r and then ϕ to find an asymptotic approximation $x_0(t,\epsilon)$ to the true solution $x(t)$ with properties

$$||x(t) - x_0(\epsilon, t)|| \leqslant C_1\epsilon \quad \text{for} \quad 0 \leqslant t \leqslant \frac{C_2}{\epsilon}, \tag{3.6}$$

where C_1 and C_2 are positive constants [21]. In other words, the method of averaging yields an asymptotically valid approximation for $x(t)$ whose error is bounded within some positive constant C_1 times the size of the perturbation over a time proportional to $1/\epsilon$. Thus it is easily seen that minimizing the size of the perturbation yields an approximation with both lower error and a longer duration of validity.

In an effort to reduce the size of the perturbing terms, several authors have investigated averaging (and other) methods that incorporate the exact solution, or a closer approximation to the exact solution, of the free vibrations of the strongly nonlinear oscillator under study. By selecting an exact solution to the nonlinear free vibrations, these authors have demonstrated improved accuracy in their results over conventional KBM averaging. It is difficult to ascertain from where this idea originated, but there are numerous examples of the approach in the literature. For example, Cap [27] and Pocobelli [28] study variations of the Langevin equation (mathematically identical to the pendulum equation) with elliptic functions and the method of averaging, and the problem of passage through resonance in the pendulum equations using elliptic functions is briefly considered in Sanders, Verhulst and Murdock's treatise on averaging and normal forms [21]. Roy studies several strongly nonlinear systems

in a series of papers [29] using elliptic functions and the method of averaging and develops a theory for the study of resonant excitation and stochastic excitation in these systems using the method of averaging. Okabe and Kondou proposed a theoretical basis for a generalized averaging method in ref. [30] and studied the Duffing oscillator in ref. [31]. Yuste, Bejarano, Garcia-Margallo and collaborators have contributed a series of articles ranging from an elliptic harmonic balance method [32, 6] to studies of nonlinear oscillators with elliptic functions as generating solutions for Krylov-Bogoliubov-type averaging methods [33, 34]. Noting that the method can be quite technical and difficult to carry out by hand, Coppola and Rand [35] developed a procedure using the computer algebra system MACSYMA to implement the method on a strongly nonlinear system. Outside of the method of averaging, Lakrad and Belhaq [36] have investigated multiple time scale methods with elliptic functions.

This investigation is organized as follows: First, a description of the system to be studied and its equations of motion are presented. In section 3.3, two averaging methods are applied to study the behavior of the system: a method using circular generating functions and a Taylor-expanded restoring force, and a method using elliptic generating functions and the fully nonlinear restoring force. Owing to the similarity of approach and the consistency of results obtained, the second method will be referred to as "Roy's Method". Then, in section 3.4, the methods are compared graphically and additional results obtained from the anharmonic method are shown to highlight system behaviors. Finally, conclusions are drawn and recommendations for future work are presented.

3.2 System description and equation of motion

The torque-forced pendulum is a cornerstone problem in dynamics and vibrations and has been reported on in a number of investigations. In this section, specifics regarding the system are rehashed and parameters are defined, and a dimensionless

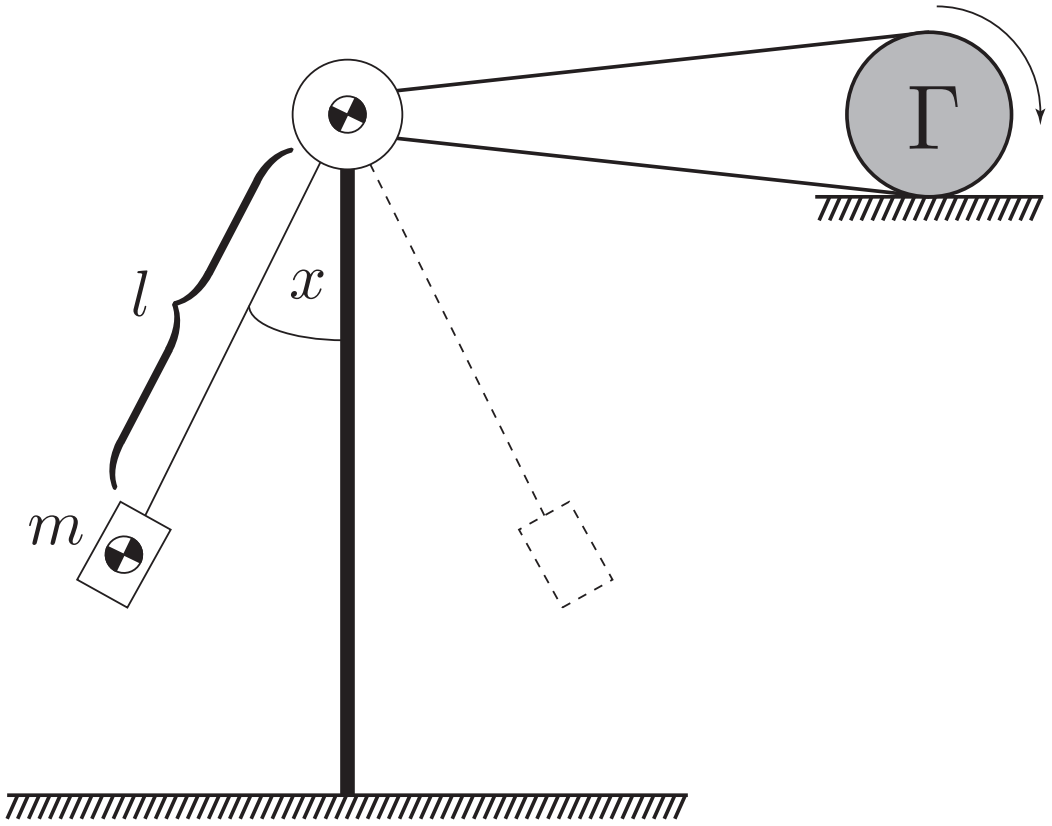

FIGURE 3.1: Schematic of a torque-excited vertical pendulum.

equation of motion is derived.

This investigation considers a pendulum mounted by a bearing onto a vertical rod, as depicted in figure 3.1. The pendulum is free to swing in one direction, denoted by the angle x, and is connected to the bearing with a rigid, massless rod of length l. A point mass m is affixed to the end of the rod and the pendulum is driven by an external harmonic torque τ_{ext}. The bearing has a viscous damping coefficient b. Given this information, the kinetic and potential energies of the system can be

written as

$$\mathcal{T} = \frac{1}{2} m l^2 (x')^2 \tag{3.7a}$$

$$\mathcal{U} = mgl(1 - \cos x), \tag{3.7b}$$

respectively, and the external nonconservative forces can be written as

$$D_x = -bx' \quad \text{(dissipative)} \tag{3.8a}$$

$$F_x = \Gamma \cos \Omega t \quad \text{(forcing)}. \tag{3.8b}$$

By applying Lagrange's equation,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathcal{T}}{\partial x'}\right) + \frac{\partial \mathcal{U}}{\partial x} = D_x + F_x, \tag{3.9}$$

the equation of motion for a harmonically forced pendulum with viscous damping,

$$ml^2 x'' + bx' + mgl \sin x = \Gamma \cos \Omega t, \tag{3.10}$$

is readily found.

It is advantageous to reduce the number of terms in the equation by rescaling time. By choosing a timescale $T = \sqrt{g/l}$ such that $t = T\tau$, a new set of parameters

$$\mu = \frac{b}{mg^{1/2} l^{3/2}}, \quad \gamma = \frac{\Gamma}{mgl}, \quad \eta = \frac{\Omega l^{1/2}}{g^{1/2}} \tag{3.11}$$

can be found yielding a rescaled form of the equation given by

$$\ddot{x} + \mu \dot{x} + \sin x = \gamma \cos \eta \tau. \tag{3.12}$$

The overdot ($\dot{}$) indicates a derivative with respect to scaled time τ. For the purposes of this investigation, we will assume that both γ and μ are small ($\ll 1$) and introduce a bookkeeping parameter ϵ to indicate their smallness, yielding

$$\ddot{x} + \sin x = \epsilon(-\mu\dot{x} + \gamma\cos\eta\tau). \tag{3.13}$$

In state-space form, the equation can be written

$$\begin{pmatrix}\dot{x}_1\\ \dot{x}_2\end{pmatrix} = \begin{pmatrix} x_2 \\ -\sin x_1 + \epsilon(-\mu x_2 + \gamma\cos\eta\tau)\end{pmatrix}. \tag{3.14}$$

Note that this vector field is rapidly varying in both x_1 and x_2. The objective of the following section will be to find a transformation that yields a slowly varying vector field to replace the original $(\dot{x}_1, \dot{x}_2)$ vector field.

3.3 Periodic solutions by averaging

The method of averaging is a well-established procedure for approximating slowly-varying vector fields. By averaging over the period(s) of the fast variation(s), a close approximation of the original vector field can be found. The purpose of averaging is to arrive at a vector field that is simple enough to admit closed form solutions and thus is amenable to analysis. For an extremely detailed overview of averaging methods and a thorough discussion of its mathematical underpinnings, the reader is encouraged to consult ref. [21].

In forced vibrations, the method of averaging is a commonly used tool for finding periodic orbits of a system, typically by transforming the state space $(\dot{x}, \ddot{x})$ into a coordinate system of amplitude- and phase-like variables (e.g. r and ϕ). Let Ω be the driving frequency of external harmonic forcing and ω be the resonant frequency of the unforced system. Then what is sought is a near-identity transformation of the

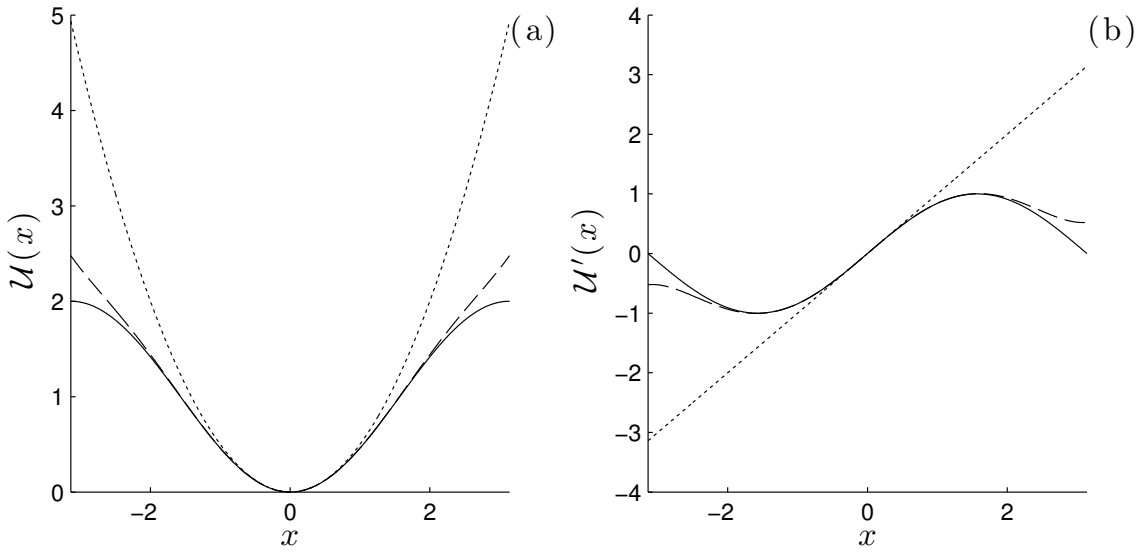

FIGURE 3.2: Potential $\mathcal{U}(x)$ (a) and restoring force $\mathcal{U}'(x)$ (b) for the vertical pendulum. The fully nonlinear potential and restoring force are shown by the solid black line, a fifth-order approximation is shown by a dashed line and the harmonic approximation is shown by a dotted line. Note that the fully nonlinear potential and restoring force are periodic and repeat from the unit cell shown above along the x axis.

form

$$\mathcal{T} : \begin{pmatrix} \dot{x} \\ \ddot{x} \end{pmatrix} \mapsto \begin{pmatrix} \dot{r} \\ \dot{\phi} \end{pmatrix} \tag{3.15}$$

with the objective of transforming the rapidly-varying nonautonomous vector field

$$\begin{pmatrix} \dot{x} \\ \ddot{x} \end{pmatrix} = \begin{pmatrix} g_1(x, \dot{x}, t) \\ g_2(x, \dot{x}, t) \end{pmatrix} \tag{3.16}$$

to a slowly-varying autonomous vector field

$$\begin{pmatrix} \dot{r} \\ \dot{\phi} \end{pmatrix} = \begin{pmatrix} 0 \\ \delta \end{pmatrix} + \epsilon \begin{pmatrix} f_1(r, \phi) \\ f_2(r, \phi) \end{pmatrix}, \tag{3.17}$$

where $\delta = \omega^2 - \Omega^2 = \mathcal{O}(\epsilon)$ is a small detuning parameter. The intermediate step that transforms the equations of state into a slowly-varying vector field of these amplitude- and phase-like variables is accomplished by assuming a particular generating solution,

so called because it "generates" the transformation[2]. The generating solution is typically an exact solution to the unperturbed equations of motion; for example, a harmonic solution is often chosen as it is usually convenient to recast a nonlinear system as a perturbed harmonic oscillator. After applying the transformation, any rapidly varying angles that arise are eliminated by averaging over their periods to ensure the system satisfies the form of equation (3.17). It can be shown that the approximation error over one period of oscillations T is bounded by a constant C_1 times the size of the perturbations ϵ as

$$\int_0^T ||x_e(t) - x_a(t)|| \, \mathrm{d}t \leqslant C_1 \epsilon, \tag{3.18}$$

where $x_e(t)$ is the exact periodic orbit of the system and $x_a(t)$ is the periodic orbit found by the averaging method, given some reasonable constraints on the smoothness and continuity of the vector fields above (see ref. [21]).

As will be shown in this investigation, the choice of generating solution determines both the depth and accuracy of analysis that can be performed. In the following subsections, two different generating solutions are utilized to transform equation (3.13) to the form of equation (3.17). Vastly different results are obtained when using each method. The two solutions are developed in the following subsections and compared in the next section.

3.3.1 Averaged vector field obtained via a harmonic generating solution

Recasting the equation of motion of a nonlinear oscillator as a perturbed simple harmonic oscillator is a common approach in the study of nonlinear vibrations. Doing so allows the investigator to assume a harmonic generating solution and incorporate the perturbing terms, including the nonlinear restoring force terms and possibly the

[2] Note that this "generating solution" has a different (less nuanced) meaning than the generating functions found in the physics literature.

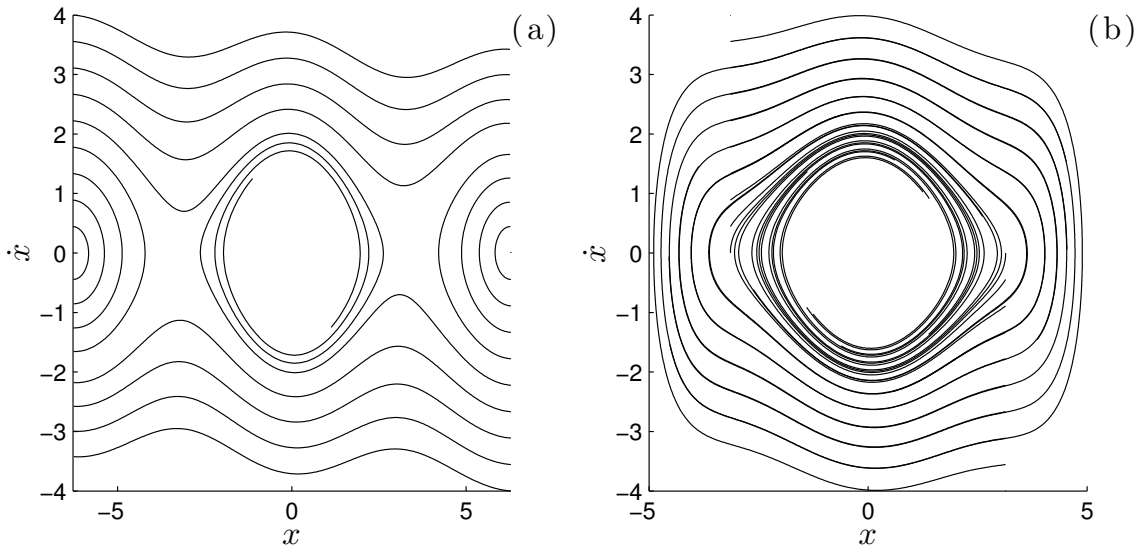

FIGURE 3.3: Comparison of global behavior for the fully nonlinear (a) and Taylor-series expanded (b) vector fields for equation (3.13). The Taylor expansion retains terms up to fifth order in the restoring force. A collection of trajectories are shown for the unforced case for each vector field. The vector fields are very similar in the neighborhood of the downward static equilibrium $(x, \dot{x}) = (0, 0)$, but differ greatly away from this neighborhood. Parameters are $\mu = 0.05$ and $\gamma = 0$.

dissipative and forcing terms, into the solution parameters. In the case of the pendulum equation, averaging the vector field obtained by assuming a harmonic generating solution yields a solution that is valid in the neighborhood of the downward equilibrium. This solution is typically accurate for light forcing and damping provided that the amplitude of the solution is moderate.

To obtain this solution, the restoring force is expanded in a Taylor series about the origin and terms up to fifth order are retained. The results of this expansion can be seen in figure 3.2 as the dashed line. Then, the nonlinear and nonconservative terms are considered small perturbations and the system is recast as a perturbed harmonic oscillator. This yields

$$\ddot{x} + x = \epsilon \left(\frac{1}{6} x^3 - \frac{1}{120} x^5 - \mu \dot{x} + \gamma \cos \eta \tau \right) \tag{3.19}$$

as the new equation of motion for the system. A comparison of the truncated vector field given by equation (3.19) to the fully nonlinear vector field can be seen in figure 3.3. Although the solution to the unperturbed system is simply $x = r\sin(t+\phi)$, where r and ϕ would be determined by the initial conditions, some forethought leads to the choice of

$$x = r\sin(\eta\tau + \phi) \tag{3.20a}$$

$$\dot{x} = r\eta\cos(\eta\tau + \phi) \tag{3.20b}$$

as a generating solution because of the frequency entrainment of periodic solutions that will occur due to the harmonic forcing. This solution is then used to transform the equations into an (r, ϕ) vector field, which is slowly varying. Let $\theta = \eta\tau + \phi$. Then, to perform the transformation, a transformation matrix

$$\boldsymbol{A} = \begin{pmatrix} \partial_r x & \partial_\theta x \\ \partial_r \dot{x} & \partial_\theta \dot{x} \end{pmatrix} = \begin{pmatrix} \sin\theta & r\cos\theta \\ \eta\cos\theta & -r\eta\sin\theta \end{pmatrix} \tag{3.21}$$

is constructed such that

$$\begin{pmatrix} \dot{r} \\ \dot{\theta} \end{pmatrix} = \boldsymbol{A}^{-1}\begin{pmatrix} \dot{x} \\ \ddot{x} \end{pmatrix} = \boldsymbol{A}^{-1}\begin{pmatrix} r\eta\cos\theta \\ -r\sin\theta + \epsilon f(x, \dot{x}, \tau) \end{pmatrix} \tag{3.22}$$

where $\epsilon f(x, \dot{x}, t) = \epsilon\left(\frac{1}{6}x^3 - \frac{1}{120}x^5 - \mu\dot{x} + \gamma\cos\eta\tau\right)$. Recalling that $\theta = \eta\tau + \phi$ and hence $\dot{\phi} = \dot{\theta} - \eta$, equation 3.22 can be rewritten as

$$\begin{pmatrix} \dot{r} \\ \dot{\phi} \end{pmatrix} = \begin{pmatrix} 0 \\ -\eta \end{pmatrix} + \boldsymbol{A}^{-1}\begin{pmatrix} r\eta\cos\theta \\ -r\sin\theta + \epsilon f(x, \dot{x}, \tau) \end{pmatrix}, \tag{3.23}$$

yielding lengthy expressions for $\dot{r}$ and $\dot{\phi}$. By expanding equation (3.23), making the substitution $\phi = \theta - \eta t$ and averaging over all non-resonant angles (i.e. θ, $n\theta$, etc.)

we find the averaged equations for r and ϕ as

$$2\eta\langle\dot{r}\rangle = \epsilon\left(\gamma\cos\phi - \mu\eta r\right) \tag{3.24a}$$

$$2\eta\langle\dot{\phi}\rangle = 1 - \eta^2 + \epsilon\left(-\frac{1}{8}r^2 + \frac{1}{192}r^4 - \frac{\gamma\sin\phi}{r}\right) \tag{3.24b}$$

Equation (3.24) is of the form given in equation (3.17); the detuning term $1-\eta^2$ vanishes near resonance and is $\mathcal{O}(\epsilon)$ within a small neighborhood of the resonant frequency of equation (3.19) ($\omega = 1$). The equilibrium solutions (r_e, ϕ_e) of equations (3.24) specify the periodic orbits of (3.13) as approximated by the steps above. Then, by recalling once more that $\theta = \eta t + \phi$, the approximate periodic orbits are given by

$$x = r_e\sin(\eta t + \phi_e) \tag{3.25a}$$

$$\dot{x} = r_e\eta\cos(\eta t + \phi_e). \tag{3.25b}$$

3.3.2 Averaged vector field obtained via an anharmonic generating solution (Roy's Method)

An alternative approach is to consider the solutions of equation (3.13) without altering the restoring force. This is the basis of the methods employed by Roy [29], Cap [27] and doubtlessly others in seeking a more accurate approximation of the orbits of equation (3.13) at high amplitudes or in the rotating domain. The unperturbed system

$$\ddot{x} + \sin x = 0 \tag{3.26}$$

has solutions

$$x_l = 2\arcsin\left[k_l \operatorname{sn}(u, k_l)\right] \tag{3.27a}$$

$$\dot{x}_l = 2k_l \operatorname{cn}(u, k_l) \tag{3.27b}$$

for the librating (i.e. oscillating, non-rotating) case and

$$x_r = 2\operatorname{am}(u/k_r, k_r) \tag{3.28a}$$

$$\dot{x}_r = \pm\frac{2}{k_r}\operatorname{dn}(u/k_r, k_r) \tag{3.28b}$$

for the rotating case, where $\operatorname{sn}(u,k)$, $\operatorname{cn}(u,k)$, $\operatorname{dn}(u,k)$ and $\operatorname{am}(u,k)$ are the Jacobi elliptic functions "sine amplitude", "cosine amplitude", "delta amplitude" and "amplitude", respectively. These functions are sometimes referred to as "anharmonic" functions; they share many properties with harmonic functions but have many unique properties of their own. Consult reference [37] for an engineering and physics oriented treatment of elliptic functions and [38] for a detailed mathematical background on the theory and application of these functions.

The solutions shown in equations (3.27) and (3.28) can be found by direct integration of equation (3.26), see reference [39] for a detailed approach. The Jacobi elliptic functions (JEFs) accept two arguments: a phase, here denoted as u, and a modulus, here denoted as k. The modulus k is restricted to the range $k \in [0, 1]$; at $k = 0$, the JEFs degenerate into trigonometric and constant functions, while at $k = 1$ the JEFs become hyperbolic functions. For equation (3.26), the heteroclinic orbits are a degenerate case of both equations (3.27) and (3.28) as $k \to 1$; in fact, the solutions are equivalent to one another when $k > 1$, which can be shown by performing a modulus transformation (see [38]).

In essence, the modulus k serves as the amplitude parameter for these solutions: as k is varied, the amplitude and shape of the solutions change. However, as k is varied, the period of the solution varies as well. This allows the amplitude dependence of solution periods to be captured naturally. While this introduces a new set of challenges in the analysis process, the goal of reducing equation (3.13) to the form of equation (3.17) is fundamentally the same.

General approach

Much of the content of this section is based on the article by V. Roy [29]. While his article as a whole considers several nonlinear systems, the section focusing on the pendulum served as the inspiration for this investigation. A key insight offered by Roy is that if a vector field can be made to be 2π-periodic in its angular variables, replacing anharmonic functions such as the JEFs with their Fourier series approximations greatly simplifies the averaging process in two ways: first, the static average is a natural byproduct of a Fourier series expansion; second, by making an appropriate choice of angle variables, the Fourier series can be used to compute the projection of an arbitrary periodic vector field on a harmonic forcing function. This eliminates the need to evaluate integrals of products of harmonic and anharmonic functions that would otherwise require the use of complex analysis. Thus, once the system is placed into the form of equation (3.17), the right-hand side will need to be Fourier expanded so that the vector field is amenable to analysis (see [29, 21]). The computer algebra system (CAS) Maple [40] is used to aid in the manipulation of the expressions that follow.

To begin, a transformation matrix

$$\boldsymbol{B} = \begin{pmatrix} \partial_k x & \partial_u x \\ \partial_k \dot{x} & \partial_u \dot{x} \end{pmatrix} \tag{3.29}$$

is obtained for both the librating ($\boldsymbol{B}_l$) and rotating solutions ($\boldsymbol{B}_r$). In practice, when using a computer algebra system (CAS), it is better to leave the terms of $\boldsymbol{B}$ unevaluated as many may cancel easily when grouped but prove challenging for the CAS to simplify when expanded. Then

$$\begin{pmatrix} \dot{k} \\ \dot{u} \end{pmatrix} = \boldsymbol{B}^{-1} \begin{pmatrix} \dot{x} \\ \ddot{x} \end{pmatrix} = \boldsymbol{B}_l^{-1} \begin{pmatrix} \dot{x} \\ -\sin(x) + \epsilon f(x, \dot{x}, \tau) \end{pmatrix} \tag{3.30}$$

in both cases. After expanding, a system of the form

$$\dot{k} = \epsilon \frac{1}{\Delta} \dot{x} f(x, \dot{x}, \tau) \tag{3.31a}$$

$$\dot{u} = 1 - \epsilon \frac{1}{\Delta} \frac{\partial x}{\partial k} f(x, \dot{x}, \tau) \tag{3.31b}$$

is found for each generating solution, where Δ is the determinant of $\boldsymbol{B}$. As noted by Roy in reference [29], the vector field in equation (3.31) is not 2π-periodic in u; thus the equation for $\dot{u}$ must be transformed to a new angular variable that is 2π-periodic by incorporating the period T of the solution. To do so, a new angular coordinate θ is defined such that $u = T\theta$. Note that T is a function of k, which is a function of τ; this yields

$$\dot{k} = \epsilon \frac{1}{\Delta} \dot{x} f(x, \dot{x}, \tau) \tag{3.32a}$$

$$\dot{u} = T\dot{\theta} - \dot{T}\theta = 1 - \epsilon \frac{1}{\Delta} \frac{\partial x}{\partial k} f(x, \dot{x}, \tau) \tag{3.32b}$$

and therefore

$$\dot{k} = \epsilon \frac{1}{\Delta} \dot{x} f(x, \dot{x}, \tau) \tag{3.33a}$$

$$\begin{aligned} \dot{\theta} &= \frac{1}{T}\left[1 + \frac{\partial T}{\partial k}\dot{k}\theta - \epsilon \frac{1}{\Delta}\frac{\partial x}{\partial k} f(x, \dot{x}, \tau)\right] \\ &= \frac{1}{T}\left[1 - \epsilon \frac{1}{\Delta}\left(\frac{\partial x}{\partial k} - \dot{x}\frac{\partial T}{\partial k}\theta\right) f(x, \dot{x}, \tau)\right]. \end{aligned} \tag{3.33b}$$

Terms in $\left(\frac{\partial x}{\partial k} - \dot{x}\frac{\partial T}{\partial k}\theta\right)$ cancel well when unexpanded, so the reader is again urged to avoid prematurely expanding or simplifying the previous equations when using a CAS.

Librating solution

For the librating solution, x and $\dot{x}$ are given by equations (3.27), $\Delta = 4k$ and $T = 4\,\mathbf{K}(k)$ where $\mathbf{K}(k)$ is the complete elliptic integral of the first kind. Thus

$$\dot{k} = \epsilon \frac{1}{2} \operatorname{cn}(u, k) f(x, \dot{x}, t) \tag{3.34a}$$

$$\dot{u} = 1 - \epsilon \frac{1}{4k}\frac{\partial x}{\partial k} f(x, \dot{x}, t). \tag{3.34b}$$

With $u = T\theta = 4\,\mathbf{K}(k)\theta$ and simplification of

$$\dot{k} = \epsilon \frac{1}{2} \operatorname{cn}(u, k) f(x, \dot{x}, t) \tag{3.35a}$$

$$\dot{\theta} = \frac{1}{4\,\mathbf{K}}\left[1 - \epsilon \frac{1}{4k}\left(\frac{\partial x}{\partial k} - 8k \operatorname{cn}(u, k)\frac{\partial\,\mathbf{K}}{\partial k}\theta\right) f(x, \dot{x}, \tau)\right]. \tag{3.35b}$$

the following equations,

$$\dot{k} = -\epsilon\mu k \operatorname{cn}^2 u + \frac{1}{2}\epsilon\left(\gamma \cos\eta\tau \operatorname{cn} u\right) \tag{3.36a}$$

$$\dot{\theta} = \frac{1}{4\,\mathbf{K}} - \epsilon \frac{1}{8kk'^2\,\mathbf{K}}\left[\operatorname{sn} u \operatorname{dn} u - \mathrm{Z}(u) \operatorname{cn} u\right]\left(-2\mu k \operatorname{cn} u + \gamma \cos\eta\tau\right), \tag{3.36b}$$

are obtained, where $(k')^2 = 1 - k^2$ is the complimentary modulus. We then expand the right-hand sides of equations (3.36) as Fourier series:

$$\dot{k} = -\epsilon\mu k \sum_{n=0}^{\infty} b_n(k) \cos(4n\pi\theta) + \frac{1}{2}\epsilon \left(\gamma \cos \eta\tau \sum_{n=0}^{\infty} c_n(k) \cos((2n-1)2\pi\theta) \right) \tag{3.37a}$$

$$\dot{\theta} = \frac{1}{4\,\mathbf{K}} - \epsilon \frac{1}{8kk'^2\,\mathbf{K}} \left[\sum_{n=1}^{\infty} (\alpha_n(k) - \beta_n(k)) \sin[(2n-1)2\pi\theta] \right] \times$$

$$\left[-2\mu k \sum_{n=1}^{\infty} c_n(k) \cos[(2n-1)2\pi\theta] + \gamma \cos \eta\tau \right] \tag{3.37b}$$

The trigonometric terms can then be combined using multiplication identities, yielding

$$\dot{k} = -\epsilon\mu k \sum_{n=0}^{\infty} b_n(k) \cos(4n\pi\theta) + \frac{1}{2}\epsilon\gamma \left(\sum_{n=0}^{\infty} c_n \left[\frac{1}{2}\cos\phi_n + \frac{1}{2}\cos\psi_n \right] \right) \tag{3.38a}$$

$$\dot{\theta} = \frac{1}{4\,\mathbf{K}} - \epsilon \frac{1}{8kk'^2\,\mathbf{K}} \left[\sum_{n=1}^{\infty} (\alpha_n(k) - \beta_n(k)) \cdot \right.$$

$$\left. \left(\left[\frac{1}{2}\sin\phi_n + \frac{1}{2}\sin\psi_n \right] - 2\mu k \sum_{n=1}^{\infty} c_n(k) \cos[(2n-1)2\pi\theta] \right) \right] \tag{3.38b}$$

where $\phi_n = (2n-1)2\pi\theta - \eta\tau$ and $\psi_n = (2n-1)2\pi\theta + \eta\tau$. Averaging over all non-resonant phase angles (i.e. eliminating all angles but ϕ_1) and substituting $\dot{\phi}_1 = 2\pi\dot{\theta} - \eta$ creates an averaged vector field

$$\langle \dot{k} \rangle = -\epsilon\mu k b_0(k) + \epsilon \frac{1}{4} \Gamma c_1(k) \cos\phi_1 \tag{3.39a}$$

$$\langle \dot{\phi}_1 \rangle = \frac{\pi}{2\,\mathbf{K}} - \eta - \epsilon \frac{\pi\gamma}{8kk'^2\,\mathbf{K}} (\alpha_1(k) - \beta_1(k)) \sin\phi_1. \tag{3.39b}$$

of the form specified in equation (3.17). The stationary points of equations (3.39) determine the parameter k and phase ϕ_1 for periodic librating orbits. These parameters can then be substituted into

$$x = 2\arcsin\left[k\,\mathrm{sn}\left(\frac{4\,\mathbf{K}}{2\pi}\left(\eta\tau+\phi_1\right),k\right)\right] \tag{3.40a}$$

$$\dot{x} = 2k\,\mathrm{cn}\left(\frac{4\,\mathbf{K}}{2\pi}\left(\eta\tau+\phi_1\right),k\right) \tag{3.40b}$$

to compute the trajectory of the periodic orbits found by the averaged equations in k and ϕ. Note that equations (3.40) are equivalent to equations (3.27). The Fourier coefficients are given by

$$b_0(k) = 1-\frac{1}{k^2}\left(1-\frac{\mathbf{E}(k)}{\mathbf{K}(k)}\right) \tag{3.41a}$$

$$c_n(k) = \frac{\pi/k\,\mathbf{K}(k)}{\cosh((n-\frac{1}{2})\pi\,\mathbf{K}(k')/\,\mathbf{K}(k))} \tag{3.41b}$$

$$\alpha_n(k) = (2n-1)\frac{\pi}{2\,\mathbf{K}(k)}c_n \tag{3.41c}$$

$$\beta_n(k) = \frac{1}{2}\sum_{m=1}^{\infty} z_m\left(c_{m-n+1}-c_{m+n}\right) \tag{3.41d}$$

$$z_n(k) = \frac{\pi/\,\mathbf{K}(k)}{\sinh(n\pi\,\mathbf{K}(k')/\,\mathbf{K}(k))} \tag{3.41e}$$

where $k' = \sqrt{1-k^2}$. The coefficients are also defined in [29] and derived in [37, 38]. In practice, the sum in equation (3.41d) is evaluated until convergence, which is defined as $|\beta_n - \beta_{n-1}| < 10^{-6}$ for the purposes of this investigation.

Rotating solution

For the rotating solution, x and $\dot{x}$ are given by equations (3.28), $\Delta = -4k^{-3}$ and $T = 2k\,\mathbf{K}(k)$. Let $v = u/k$. Then

$$\dot{k} = -\epsilon\frac{1}{2}k^2\,\mathrm{dn}(v,k)f(x,\dot{x},t) \tag{3.42a}$$

$$\dot{u} = 1 + \epsilon\frac{k^3}{4}\frac{\partial x}{\partial k}f(x,\dot{x},t). \tag{3.42b}$$

With $u = T\theta = 2k\,\mathbf{K}(k)\theta$ and simplification of

$$\dot{k} = -\epsilon\frac{1}{2}k^2\,\mathrm{dn}(v,k)f(x,\dot{x},t) \tag{3.43a}$$

$$\dot{\theta} = \frac{1}{2k\,\mathbf{K}}\left[1 - \epsilon\frac{k^3}{4}\left(\frac{\partial x}{\partial k} - \frac{8}{k}\,\mathrm{dn}(v,k)\frac{\partial\,\mathbf{K}}{\partial k}\theta\right)f(x,\dot{x},\tau)\right]. \tag{3.43b}$$

the following equations,

$$\dot{k} = \epsilon\mu k\,\mathrm{dn}^2\,v - \frac{1}{2}\epsilon k^2\,(\gamma\cos\eta\tau\,\mathrm{cn}\,u) \tag{3.44a}$$

$$\dot{\theta} = \frac{1}{2k\,\mathbf{K}} - \epsilon\frac{k}{4k'^2\,\mathbf{K}}\left[k^2\,\mathrm{sn}\,v\,\mathrm{cn}\,v - \mathrm{Z}(v)\,\mathrm{dn}\,v\right)\left(-\frac{2}{k}\mu\,\mathrm{dn}\,v + \gamma\cos\eta\tau\right), \tag{3.44b}$$

are obtained, where again $(k')^2 = 1 - k^2$ is the complimentary modulus. We then

expand the right-hand sides of equations (3.44) as Fourier series:

$$\dot{k} = \epsilon\mu k \sum_{n=0}^{\infty} a_n(k) \cos(2n\pi\theta) - \frac{1}{2}\epsilon k^2 \left(\gamma \cos\eta\tau \sum_{n=0}^{\infty} d_n(k) \cos(2n\pi\theta) \right) \tag{3.45a}$$

$$\dot{\theta} = \frac{1}{2k\,\mathbf{K}} - \epsilon \frac{k}{4k'^2\,\mathbf{K}} \left[\sum_{n=1}^{\infty} \left(k^2\zeta_n(k) - \delta_n(k) \right) \sin(2n\pi\theta) \right] \times$$

$$\left[-\frac{2}{k}\mu \sum_{n=1}^{\infty} d_n(k) \cos(2n\pi\theta) + \gamma\cos\eta\tau \right] \tag{3.45b}$$

The trigonometric terms can then be combined using multiplication identities, yielding

$$\dot{k} = \epsilon\mu k \sum_{n=0}^{\infty} a_n(k) \cos(2n\pi\theta) - \frac{1}{2}\epsilon\gamma k^2 \left(\sum_{n=0}^{\infty} d_n \left[\frac{1}{2}\cos\phi_n + \frac{1}{2}\cos\psi_n \right] \right) \tag{3.46a}$$

$$\dot{\theta} = \frac{1}{2k\,\mathbf{K}} - \epsilon \frac{k}{4k'^2\,\mathbf{K}} \left[\sum_{n=1}^{\infty} \left(k^2\zeta_n(k) - \delta_n(k) \right) \cdot \right.$$

$$\left. \left(\left[\frac{1}{2}\sin\phi_n - \frac{1}{2}\sin\psi_n \right] - \frac{2}{k}\mu \sum_{n=1}^{\infty} d_n(k) \cos(2n\pi\theta) \right) \right] \tag{3.46b}$$

where $\phi_n = 2n\pi\theta - \eta\tau$ and $\psi_n = 2n\pi\theta + \eta\tau$. Averaging over all non-resonant phase angles (i.e. eliminating all angles but ϕ_1) and substituting $\dot{\phi}_1 = 2\pi\dot{\theta} - \eta$ creates an averaged vector field

$$\langle \dot{k} \rangle = \epsilon\mu k a_0(k) - \epsilon \frac{1}{4}\gamma k^2 d_1(k) \cos\phi_1 \tag{3.47a}$$

$$\langle \dot{\phi}_1 \rangle = \frac{\pi}{k\,\mathbf{K}} - \eta - \epsilon \frac{\pi\gamma k}{4k'^2\,\mathbf{K}} \left(k^2\zeta_1(k) - \delta_1(k) \right) \sin\phi_1. \tag{3.47b}$$

of the form specified in equation (3.17). The stationary points of equations (3.47) determine the parameter k and phase ϕ_1 for periodic rotating orbits. These parameters can then be substituted into

$$x = 2\,\mathrm{am}\left(\frac{2\,\mathbf{K}}{2\pi}(\eta\tau + \phi_1), k\right) \tag{3.48a}$$

$$\dot{x} = \frac{2}{k}\,\mathrm{dn}\left(\frac{2\,\mathbf{K}}{2\pi}(\eta\tau + \phi_1), k\right) \tag{3.48b}$$

to compute the trajectory of the periodic orbits found by the averaged equations in k and ϕ. Note that equations (3.48) are equivalent to equations (3.28). The Fourier coefficients are given by

$$a_0(k) = \frac{\mathbf{E}(k)}{\mathbf{K}(k)} \tag{3.49a}$$

$$d_n(k) = \frac{\pi/\,\mathbf{K}(k)}{\cosh(n\pi\,\mathbf{K}(k')/\,\mathbf{K}(k))} \tag{3.49b}$$

$$\zeta_n(k) = n\frac{\pi}{k^2\,\mathbf{K}(k)}d_n \tag{3.49c}$$

$$\delta_n(k) = \frac{1}{2}\sum_{m=1}^{\infty}\alpha_m\left(z_{m+n-1} - z_{m-n}\right) \tag{3.49d}$$

where $k' = \sqrt{1-k^2}$ and α_n and z_n are defined as above. The coefficients are also defined in [29] and derived in [37, 38]. In practice, the sum in equation (3.49d) is evaluated until convergence, which is defined as $|\delta_n - \delta_{n-1}| < 10^{-6}$ for the purposes of this investigation.

3.4 Results

In this section, the averaging methods described above are implemented to predict the existence and characteristics of period-1 orbits in equation (3.13). Results ob-

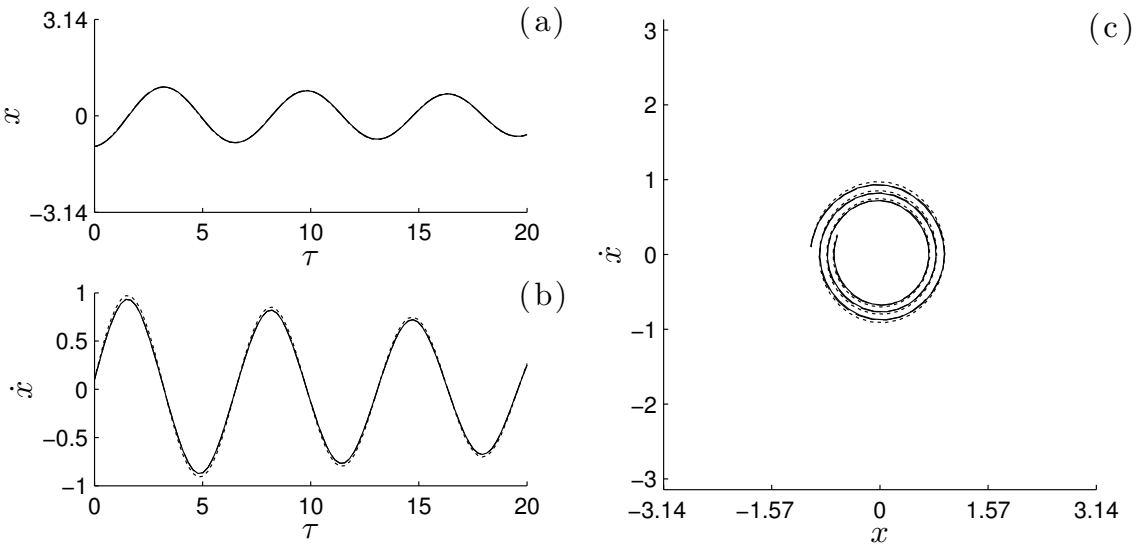

FIGURE 3.4: Comparison of amplitude decay predictions found by the anharmonic (dashed lines) and harmonic (dotted lines) averaging methods against numerical simulation (solid lines) of equation (3.13) for $\mu = 0.04$ and $(x_0, \dot{x}_0) = (-1.0, 0.1)$. Shown are the angle (a), angular velocity (b) and phase portrait (c) of the response.

tained from the anharmonic and harmonic generating solutions are compared where appropriate comparisons may be drawn, and used to highlight interesting behaviors of the system in the course of the comparisons. Unique characteristics of the averaged vector field derived from the anharmonic generating solution are also presented, allowing some insights into the rotating behaviors of the equation (3.13). Within this section, the first subsection focuses on comparisons between the predicted amplitude decay and period-1 librating orbits obtained from assuming harmonic and anharmonic generating solutions. The similarities and differences in solutions obtained by the two methods are highlighted and conclusions are drawn regarding the efficacy of each approach. The second subsection focuses solely on the period-1 rotating solution that is only obtainable from the anharmonic generating solution.

As neither the anharmonic generating solution nor the harmonic generating solution transforms equation (3.13) into averaged vector field that admits a closed-form solution for its equilibria, numerical root-finding methods are used to determine the

equilibria (k_e, ϕ_e) and (r_e, ϕ_e) for the anharmonic and harmonic averaged vector fields, respectively. To determine the stability of predicted orbits, a numerical linearized stability analysis was conducted about the equilibrium points found in the root-finding routine. The Jacobian matrix $\boldsymbol{J}$ at the equilibrium is a natural byproduct of a Newton-like root-finding algorithm, so stability could be checked with minimal additional computational cost. Orbits are considered stable when all eigenvalues λ_i of $\boldsymbol{J}$ have negative real parts (i.e. $\mathfrak{Re}(\lambda_i) < 0 \, \forall i$).

To conduct parameter studies, numerical continuation techniques were used to follow solution branches. One-parameter parameter studies follow the equilibria of the averaged vector fields as the parameter of interest is varied; a psuedo-arclength criterion is added to ensure that the system of equations remain square. Two-parameter studies follow the equilibria at bifurcation points given by $\det(\boldsymbol{J}) = \prod_{i=1}^{dim(\boldsymbol{J})} \lambda_i = 0$, adding an additional degree of freedom that allows a second parameter to be varied while keeping the system of equations square.

3.4.1 Comparison of harmonic and anharmonic approaches

In this subsection, the results obtained by assuming harmonic and anharmonic generating solutions are compared. As a first point of comparison, the decaying response under no forcing ($\gamma = 0$) is considered. Figure 3.4 shows the high accuracy of both solutions for a low-amplitude initial condition $(x_0, \dot{x}_0) = (-1.0, 0.1)$. Both solutions are capable of capturing the behavior of the system, as the oscillations still maintain a harmonic character at this low amplitude. Figure 3.5 shows the loss of accuracy for an initial condition $(x_0, \dot{x}_0) = (-3.0, 0.1)$ near the separatrix. There is a noticeable phase error in both solutions, but while the anharmonic solution retains the shape (amplitude) of the numerical solution, the harmonic solution exhibits noticeable amplitude errors.

Figure 3.6 shows a comparison of both approaches against a numerically inte-

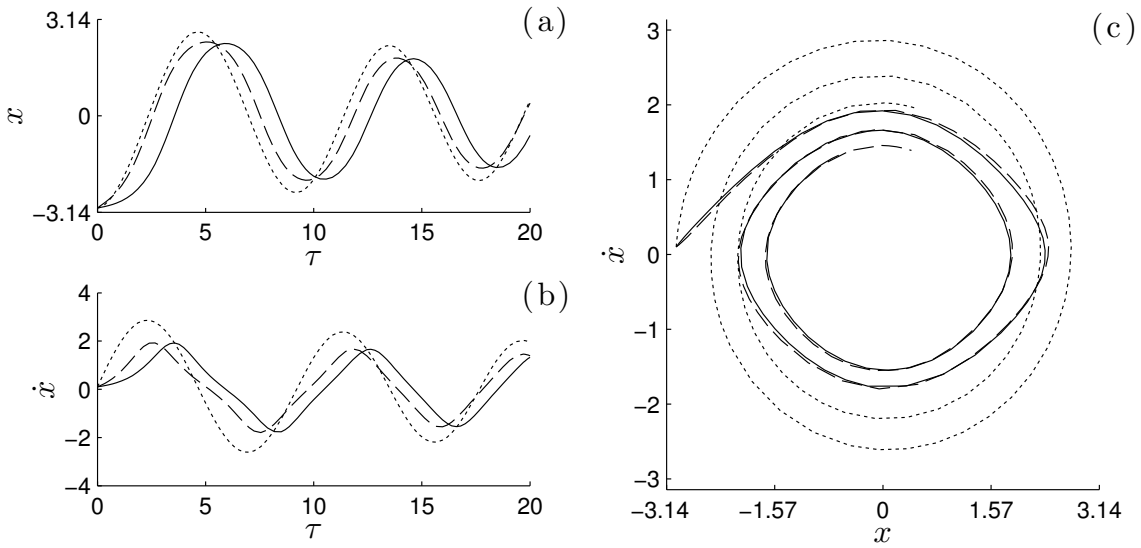

FIGURE 3.5: Comparison of amplitude decay predictions found by the anharmonic (dashed lines) and harmonic (dotted lines) averaging methods against numerical simulation (solid lines) of equation (3.13) for $\mu = 0.04$ and $(x_0, \dot{x}_0) = (-3.0, 0.1)$. Shown are the angle (a), angular velocity (b) and phase portrait (c) of the response.

grated orbit of equation (3.13). Both predicted solutions show excellent agreement with numerical simulation, although both slightly overestimate the amplitude of the solution and the harmonic generating function underestimates the peak velocities of the solution. Both solutions remain well within the predicted error tolerance of $\epsilon \approx 0.1$ as dictated by $\gamma = 0.1$.

For some parameter values, coexisting orbits exist. Figure 3.7 shows three coexisting orbits for $\gamma = 0.1$, $\mu = 0.07$ and $\eta = 0.87$. Predictions from both the harmonic and anharmonic averaged vector fields yield accurate estimates of the system behavior as determined by numerical integration of equation (3.13).

In order to find common grounds for comparison, both response amplitudes are converted to the equivalent energy determined by the potential function

$$E \equiv \mathcal{U}(x) = 1 - \cos x. \tag{3.50}$$

For the anharmonic generating function in the librating regime this corresponds to

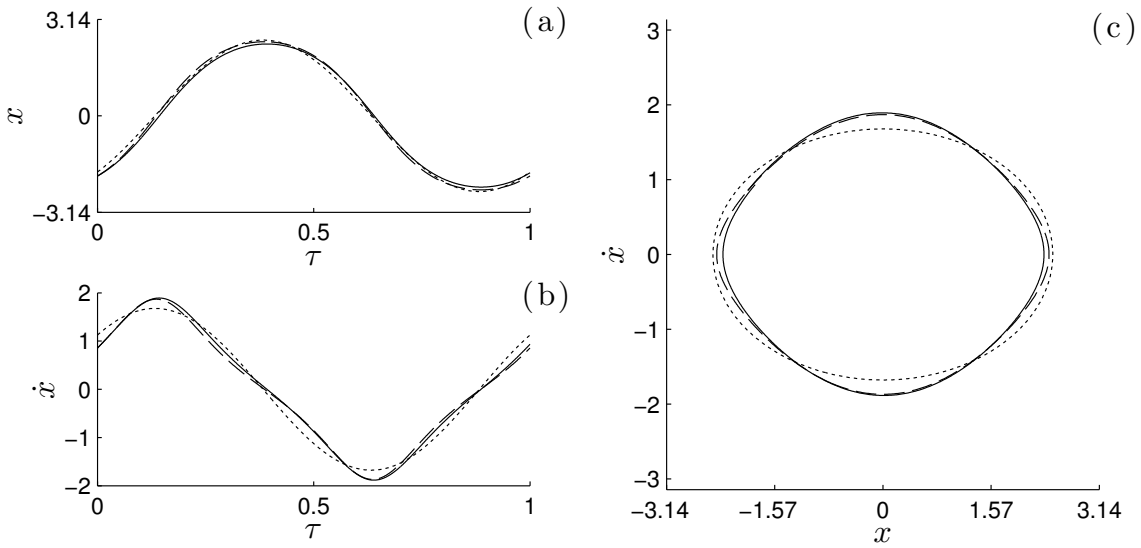

FIGURE 3.6: Comparison of a predicted stable orbit found by the anharmonic (dashed lines) and harmonic (dotted lines) averaging methods against numerical simulation (solid lines) of equation (3.13) for $\gamma = 0.1$, $\mu = 0.04$ and $\eta = 0.68$. Shown are the angle (a), angular velocity (b) and phase portrait (c) of the response.

$E = 2k_e^2$; for the harmonic generating function $E = 1 - \cos(r_e)$. By comparing the energy of the predicted solutions, disparities in scale between predicted values of k from the anharmonic generating solution and predicted values of r from the harmonic generating solution are avoided.

For a first glance into the frequency response of the system, it is useful to determine the amplitude and period of the undamped free response (i.e. $\gamma = \mu = 0$). Figure 3.8 shows the free response spines predicted by the anharmonic generating solution and the harmonic generating solution. From these spines, it is apparent that the system exhibits softening behavior. Forced response curves will exhibit their peak along the spines shown in 3.8. In figure 3.8(b), the spine predicted from the anharmonic generating solution is underlaid for comparison with the predicted free response from the harmonic generating solution. Agreement is excellent until $E \approx 1.4$, whereupon the anharmonic generating function begins to predict a slightly more pronounced softening response.

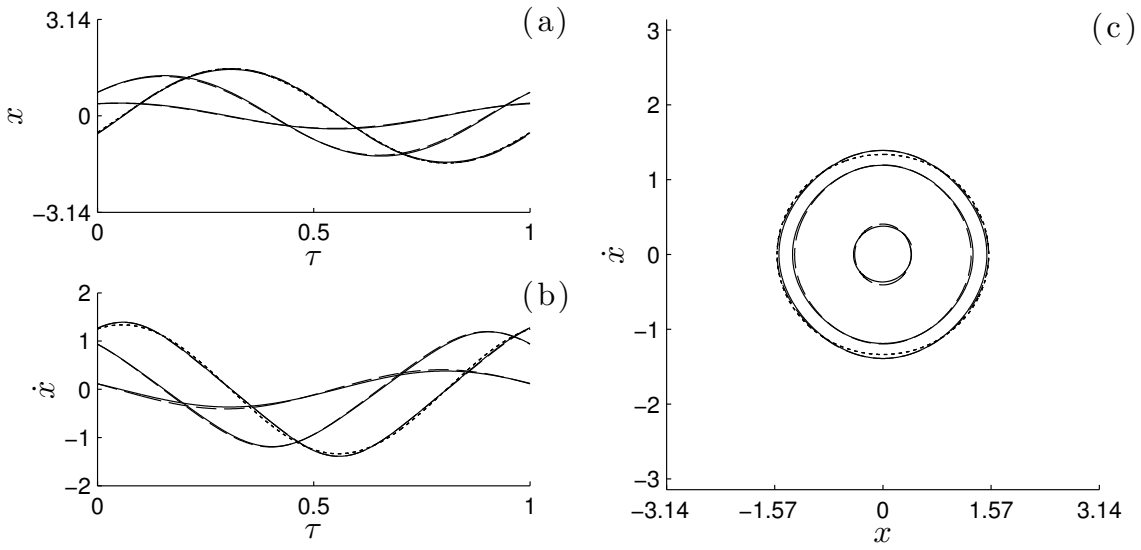

FIGURE 3.7: Comparison predicted coexisting orbits found by the anharmonic (dashed lines) and harmonic (dotted lines) averaging methods against numerical simulation (solid lines) of equation (3.13) for $\gamma = 0.2$, $\mu = 0.1$ and $\eta = 0.8$. Shown are the angle (a), angular velocity (b) and phase portrait (c) of the response.

Figure 3.9 and figure 3.10 show several frequency response curves for various values of the damping parameter μ to highlight the suppressing effects of higher damping. In both figures, it is seen that increasing the damping parameter μ while holding constant the forcing amplitude γ suppresses the domain of existence of coexisting solutions along the η axis. For sufficiently high values of μ, coexisting solutions are eliminated altogether.

Figure 3.9(a) shows the energy response as predicted by the anharmonic generating solution while figure 3.9(b) shows the energy response predicted by the harmonic generating solution; figure 3.9(c) and 3.9(d) shows phase responses of the anharmonic and harmonic generating solutions, respectively. Because of the low forcing amplitude $\gamma = 0.05$, there is excellent agreement in the predictions made by both methods in both the amplitude and phase response of the system. The pendulum has a softening response due to the shape of its potential well. The solution agreement is due to the low error in the approximation of the potential well by the Taylor series

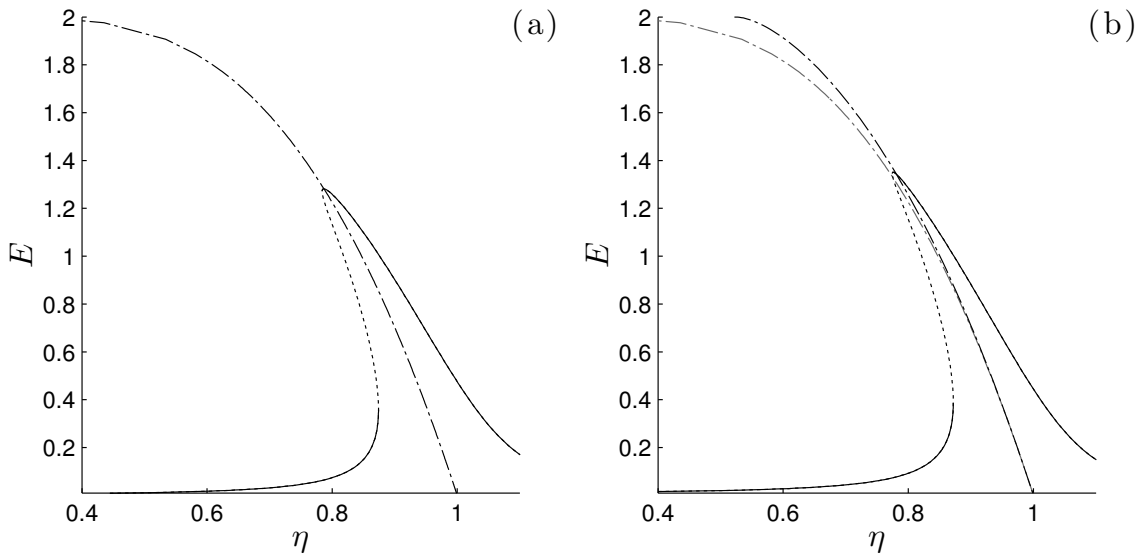

FIGURE 3.8: Comparison of the predicted resonance spines (dash-dotted lines) from the anharmonic (a) and harmonic (b) generating solutions representing the amplitude and frequency of free vibrations (i.e. $\gamma = \mu = 0$). In figure (b) the anharmonic spine is shown in light gray for comparison. Note the slightly more pronounced softening behavior predicted by the anharmonic model. These spines form a backbone for all possible frequency response curves; an example curve is shown with parameters $\gamma = 0.15$ and $\mu = 0.1$. Stable solutions exist on the solid lines while unstable solutions are found on the dotted lines.

expansion used in conjunction with the harmonic generating solution, as seen in figure 3.2. As would be expected in a nonlinear system, for reasonably low values of μ multiple coexisting solutions are found near resonance: specifically, three coexisting solutions, two stable and one unstable, are found to the left of the resonant peak.

As the forcing amplitude increases, the predictions from the anharmonic and harmonic generating solutions begin to diverge. Figure 3.10 presents energy and phase response curves that illustrate the breakdown in solution agreement as $E \to 2$. The physical reason for this breakdown is the existence of the separatrix at $E = 2$; beyond the separatrix, the pendulum system transitions into hybrid rotating-librating behaviors or purely rotating behaviors that are not predicted by either method. The predictions generated by the anharmonic method in figure 3.10(a)

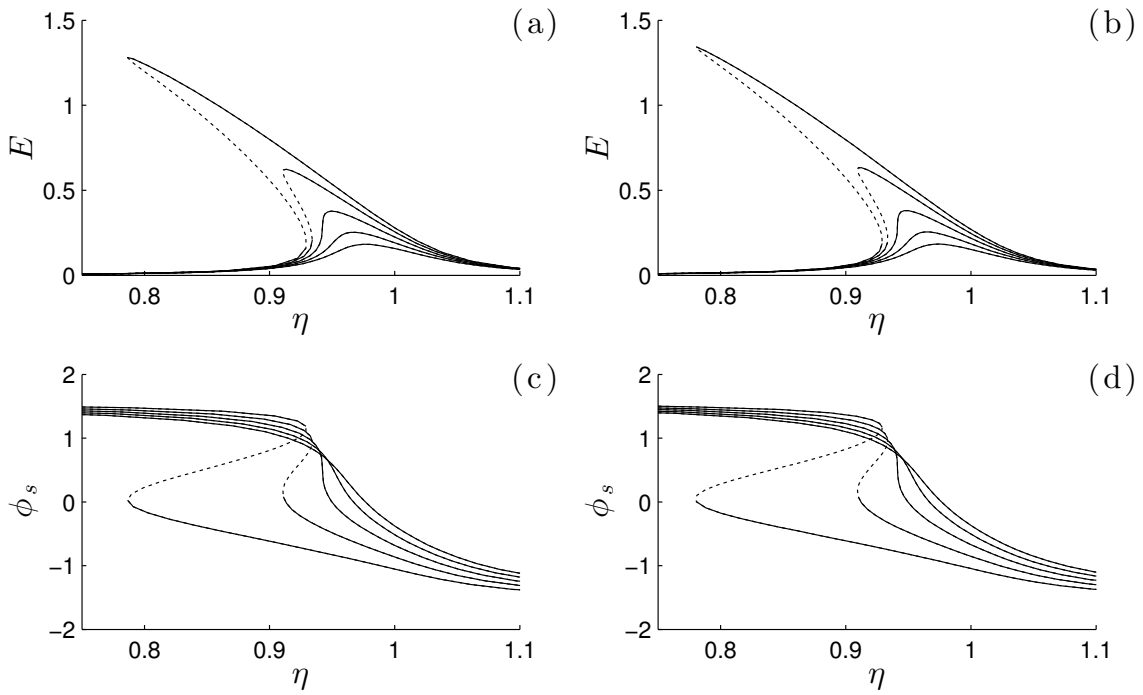

FIGURE 3.9: Comparison of predicted frequency response curves for $\mu = \{0.040, 0.055, 0.070, 0.085, 0.100\}$, $\gamma = 0.05$. Solid lines indicate stable period-1 orbits while dashed lines indicate unstable orbits. Figures (a) and (c) show the energy and phase response given by the nonlinear averaging method, while figures (b) and (d) show the energy and phase response given by the harmonic averaging method.

indicate the existence of an unstable branch of solutions at $E = 2$ beyond $\eta \approx 0.7$; these solutions are spurious and cannot be recreated by numerical integration. Indeed, as $E \to 2$, it is difficult to accurately recreate the predicted solutions with numerical simulation. Additionally, the solution curve is broken into two parts, one on each side of the resonant spine, for sufficiently high γ and low μ. All of these observations are due to complex physical behavior that occurs as the separatrix begins to break under high-amplitude forcing. In contrast, the solutions predicted by the harmonic generating function simply increase in amplitude as μ is decreased, as shown in figure 3.10(b). This is due to the lack of a separatrix in the Taylor expanded vector field. The distorted peak is due to the energy mapping from equation (3.50).

For an even more apparent view into the disparities in predictions as $E \to 2$,

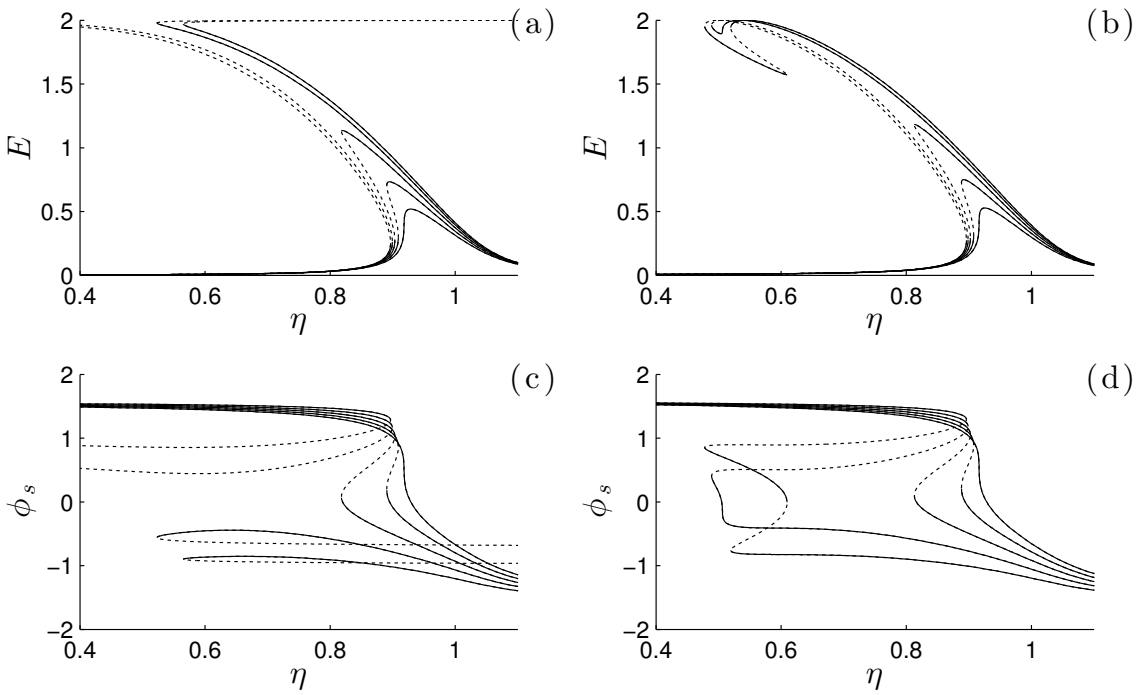

FIGURE 3.10: Comparison of predicted frequency response curves for $\mu = \{0.040, 0.055, 0.070, 0.085, 0.100\}$, $\gamma = 0.1$. Solid lines indicate stable period-1 orbits while dashed lines indicate unstable orbits. Figures (a) and (c) show the energy and phase response given by the nonlinear averaging method, while figures (b) and (d) show the energy and phase response given by the harmonic averaging method.

figure 3.11 shows the energy and phase response as the forcing is varied for $\eta = 0.8$ and several values of μ. As seen in figures 3.9 and 3.10, increasing μ while holding η constant has the effect of suppressing the domain of existence coexisting orbits along the γ axis. Coexisting high- and low-amplitude orbits for γ near 0.15 are clearly visible; however as γ increases beyond $\gamma \approx 0.35$ the anharmonic generating solution ceases to predict librating orbits and instead predicts a fourth coexisting orbit that approaches the separatrix, while the harmonic generating solution maintains predictions of librating orbits indefinitely.

As a final point of comparison between the anharmonic and harmonic generating solutions, figure 3.12 shows the domain of coexisting solutions in the (μ, γ) plane for various values of η. Inside the shaded regions, the viscous damping μ is low enough

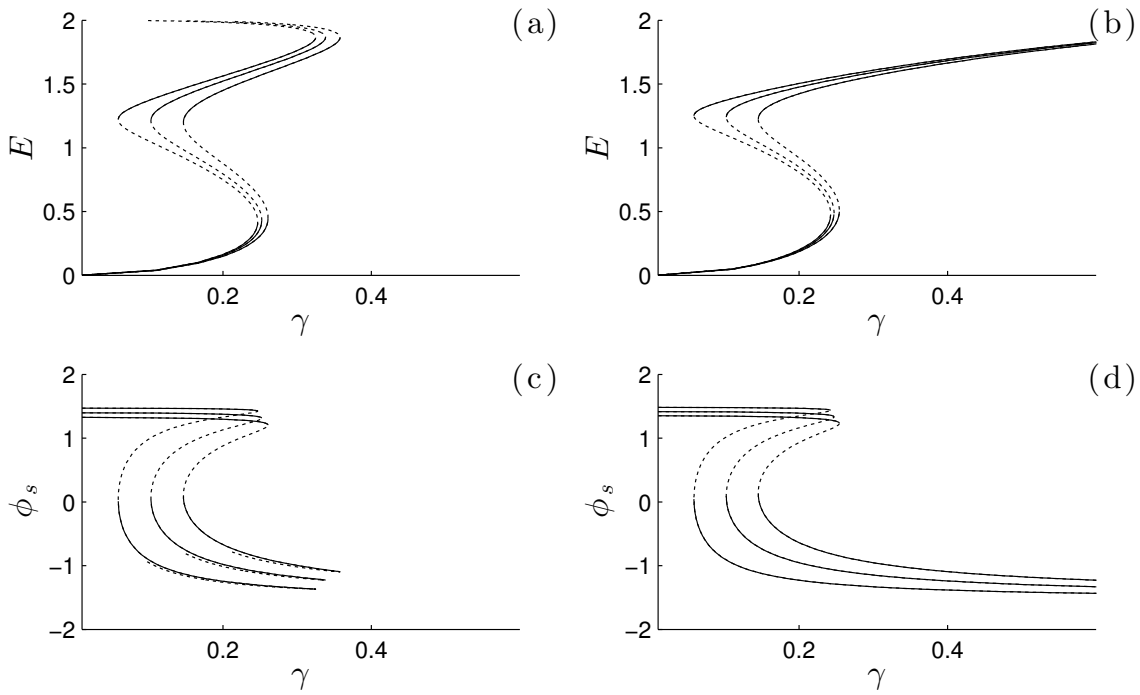

FIGURE 3.11: Forcing response for $\mu = \{0.040, 0.055, 0.070, 0.085, 0.100\}$, $\eta = 0.8$. Figures (a) and (c) show the energy and phase response given by the nonlinear averaging method, while figures (b) and (d) show the energy and phase response given by the harmonic averaging method.

and the forcing amplitude γ is high enough that solutions coexist for the forcing frequency η shown to the right of each region. As seen in several of the previous figures, the domain of existence of coexisting solutions in the η line occurs when η is slightly less than 1. In figure 3.12(b), the predictions from the anharmonic generating solution are underlaid as dotted lines for comparison. Agreement is excellent for higher values of η, but as η decreases, the higher-amplitude resonant peaks distort predictions made by the harmonic generating solution.

3.4.2 Rotating solutions from the anharmonic approach

A unique feature of the anharmonic generating solution is its ability to describe rotating orbits for $E = 2$ (corresponding to the change of modulus equation, see [38]). As rotating solutions complete one entire orbit per period, the concept of

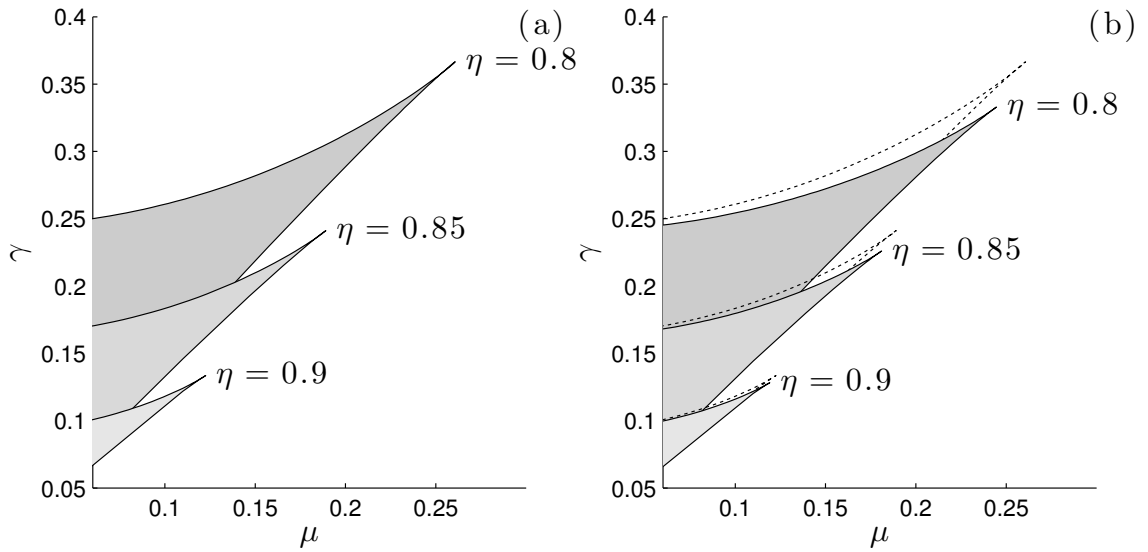

FIGURE 3.12: Domain of coexisting librating solutions in the (μ, γ) plane for $\eta = \{0.8, 0.85, 0.9\}$. Figure (a) shows the domains as estimated by the anharmonic averaging method, while figure (b) shows the domains estimated by the harmonic averaging method as solid patches with the predictions from (a) underlaid as dotted lines for reference.

amplitude is not informative in characterizing their behavior. Instead, response curves are shown with energy as defined above in equation (3.50). In the rotating case, $E = 2k^{-2}$. Rotating solutions are indeed periodic orbits given the periodicity of the rotation angle $x \in (-\pi, \pi]$ (e.g. $x = 2\pi$ is equivalent to $x = 0$); in this subsection the influence of forcing and damping on the response and existence of periodic rotating orbits is examined.

Figure 3.13 shows an example rotating orbit for $\gamma = 0.4$, $\eta = 2.5$ and $\mu = 0.01$. There is reasonable agreement between the predicted orbit and a numerically simulated orbit with the same initial conditions. The orbit shown is a stable orbit; stable rotating orbits tend only to exist when damping is low and the forcing amplitude is high. Note that the mirror image orbit also exists.

In similar fashion to figure 3.8, the undamped free rotations of the system can be studied to form a resonance spine. Figure 3.14(a) shows the resonance spine for

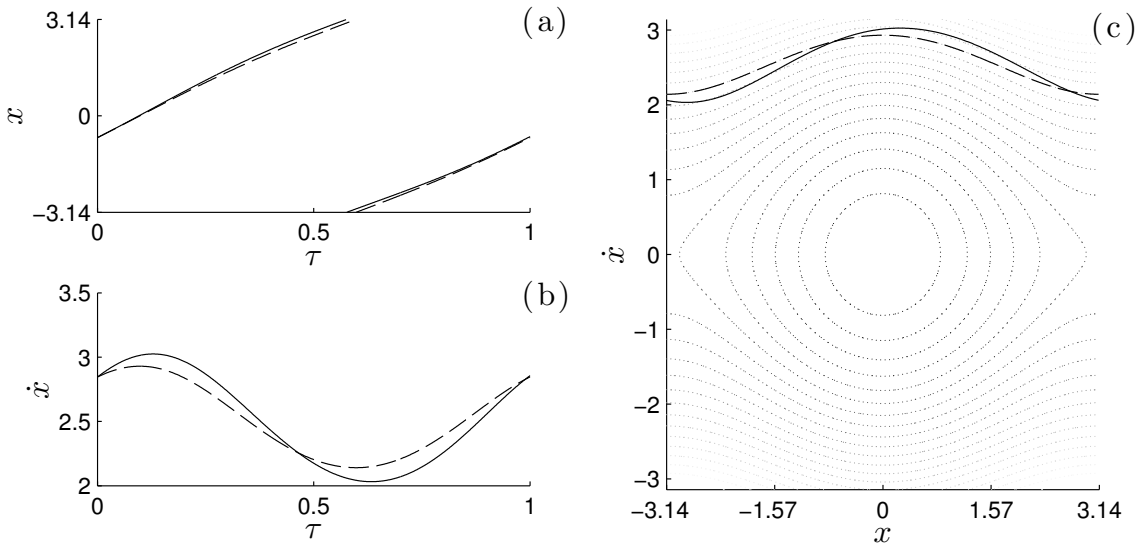

FIGURE 3.13: Rotating orbit found by the anharmonic generating solution (dashed line) compared to numerical simulation (solid line) of equation (3.13) for $\gamma = 0.4$, $\eta = 2.5$ and $\mu = 0.01$. Shown are the angle (a), angular velocity (b) and phase portrait (c) of the response. In (c), contours of the Hamiltonian, corresponding to the orbits of the unperturbed system, are shown by faint dotted lines.

both librating and rotating orbits as well as a predicted frequency response for each. Inset 3.14(b) focuses on the area where the resonance spines meet as the asymptote towards the heteroclinic orbit at $E = 2$ and $\eta = 0$. The spine of rotating orbits asymptotes to $E = 2$ quickly for $\eta < 1.5$; stable forced rotating orbits near $E = 2$ are difficult if not impossible to find.

Figure 3.15 demonstrates the effects of increasing μ on the response of rotating orbits. As seen with librating orbits, increasing μ suppresses the response energy. However, while increasing μ gradually reduced the domain of coexisting orbits in the librating case, increasing μ in the rotating case gradually eliminates the domain of stable orbits altogether.

Figure 3.16 shows the domains of existence of stable rotating orbits in the (μ, γ) plane for various values of η. Stable rotating orbits were not observed for $\eta < 1$.

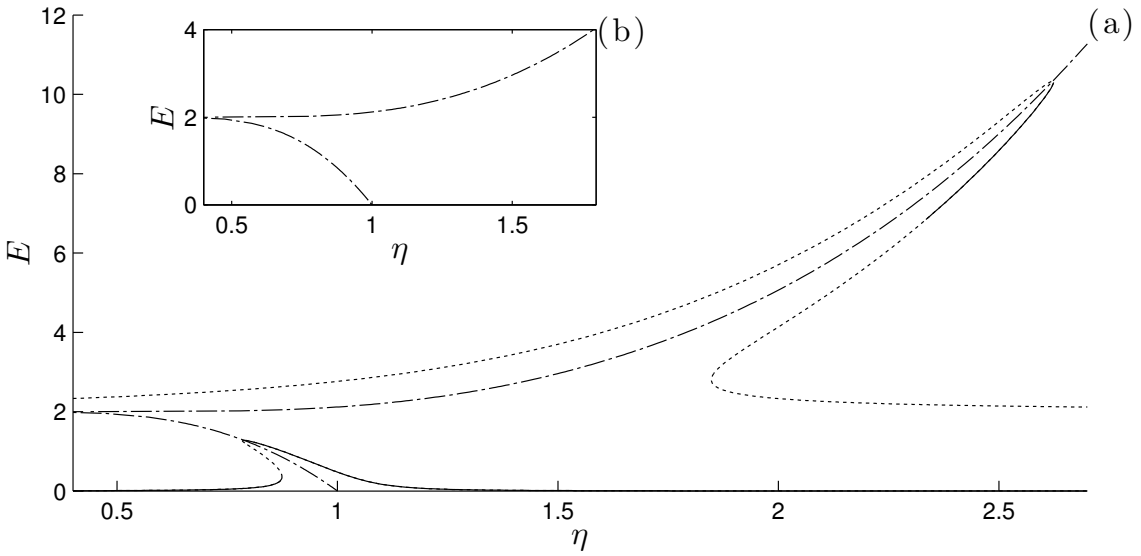

FIGURE 3.14: Resonance spines for the librating ($E < 2$) and rotating ($E > 2$) solutions as predicted using the anharmonic generating solution. Figure (a) shows both the librating and rotating spines and example frequency response curves ($\gamma = 0.15$, $\mu = 0.1$ and $\gamma = 1.1$, $\mu = 0.03$ for the librating and rotating cases, respectively. Inset (b) highlights the convergence of both spines to the heteroclinic orbit.

3.5 Conclusions

From this investigation it is simple to conclude that the appropriate choice of generating solution depends on the information one wishes to obtain from an approximate solution. The anharmonic generating solution delivered results capable of describing behaviors near the separatrix and in the rotating regime that were not well described by the harmonic generating solution. However, the effort required to obtain the averaged equations from the anharmonic generating solution is substantially higher than what is required when using a harmonic generating solution. If an investigator were solely interested in low-amplitude behaviors about the stable static equilibrium, or perhaps even designing a control scheme to stabilize the upright unstable equilibrium, an approximation that is valid in the neighborhood of the appropriate equilibrium point could prove sufficiently descriptive for the task at hand while reducing the symbolic computation workload. On the other hand, if an investigator were more

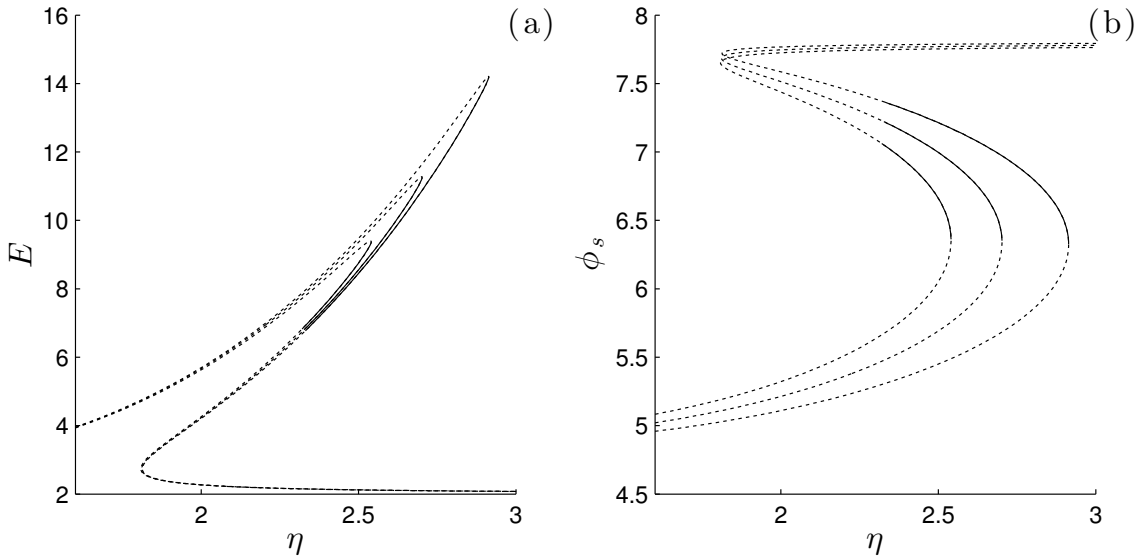

FIGURE 3.15: Predicted frequency response of rotating orbits for $\gamma = 1$, $\mu = \{0.020, 0.025, 0.030\}$. Solid lines indicate stable period-1 rotating orbits. Figure (a) shows the energy response while figure (b) shows the phase response of the orbits.

interested in studying the behavior of rotating solutions, the more descriptive anharmonic generating solution would be the better option. Further still, investigation of bifurcation behaviors such as period-doubling or subharmonic responses would require yet another approach beyond what is detailed here. Again, the investigator must use his or her discretion in choosing an appropriate generating solution so that the information sought may be obtained.

Future work in this specific area offers somewhat limited possibilities, as this investigation focuses more on comparing existing work in the literature. The field of applied mathematics has made substantial developments in the theory and application of averaging methods over the past half-century, and it is a widely adopted technique in the study of applied differential equations in fields ranging from engineering to biology. It is the author's hope that the results presented herein offer some useful perspective and intuition to the vibrations community with regards to choosing an appropriate generating solution for an analytical investigation.

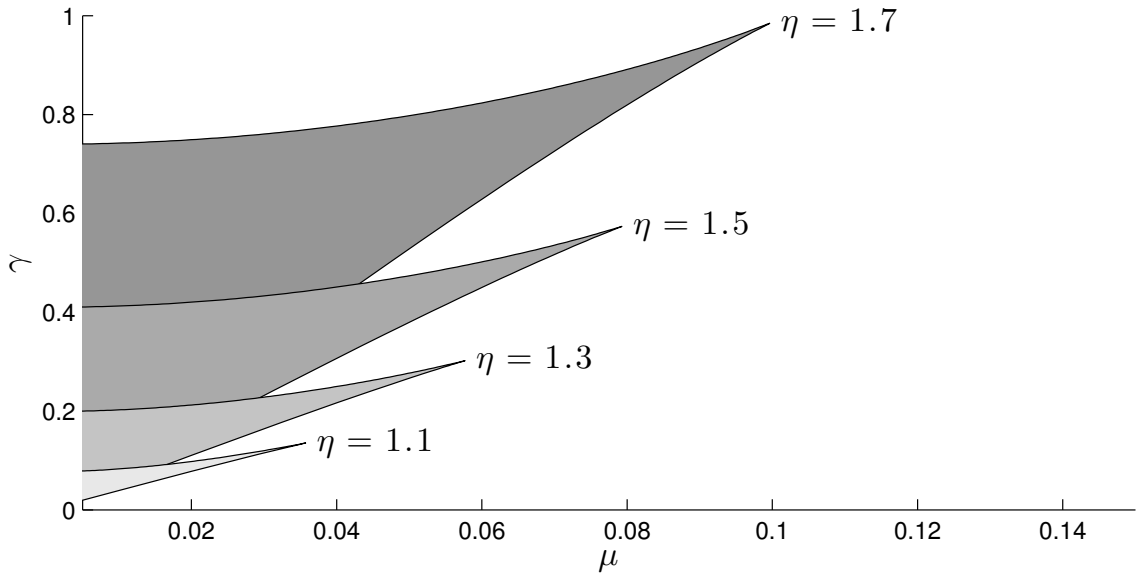

FIGURE 3.16: Domain of existence of rotating orbits in the (μ, γ) plane for $\eta = \{1.1, 1.3, 1.5, 1.7\}$. Stable rotating orbits exist in the shaded areas.

4

Numerical Bifurcation, Stability and Uncertainty Analysis of the Nonlinear Energy Harvester Equations

A common approach in the study of nonlinear vibrations is to study the periodic orbits of a system by either solving the equations exactly or approximately to arrive at a symbolic representation of the response of the system. Such an approach is typically referred to as an analytical method, and it provides a general solution for arbitrary parameter values that usually requires little computational effort to evaluate for specific values of the parameters. An example of such an approach is found in the previous chapter. Analytical methods have the advantages of generality and abstraction, but they can be difficult to obtain in strongly nonlinear systems while still maintaining a reasonable degree of accuracy and readability of the result.

A second approach is to utilize computational techniques to compute the orbits, yielding a numerical result directly. While this is certainly less general, as only one result is obtained from a computation, it is possible to use numerical techniques in an intelligent way to conduct a guided exploration of the parameter space in much

the same way as would be done with an analytical solution. By casting the problem of finding periodic orbits of a differential equation as a boundary value problem, numerical techniques such as finite difference, finite element or spectral methods can be employed to compute the orbits. Then, by using numerical continuation techniques, parameter studies can be conducted without a need for an analytical representation of the periodic orbits. This is advantageous for systems with orbits that are difficult to determine analytically or when a higher fidelity representation of model behavior is required.

In developing this methodology for the study of the system in chapter 5, it was found that the method is well suited for the analysis of equations describing nonlinear energy harvesting devices. Recent investigations conducted by S. Stanton and B. Mann led to the development of mathematical models for lead-zirconate-titanate (PZT) coated cantilever-beam energy harvester devices with nonlinear restoring forces arising from magnet-magnet interactions at the free end of the beam [4, 3]. These models have taken hold in the literature and are the standard for modeling any similar device. The magnet-magnet interactions contribute nonlinear stiffness effects that necessitate the use of approximate analytical or numerical methods in the analysis of the model equations.

The primary focus of this chapter is to adapt the models presented in [4, 3, 8, 7] into a simplified general model and analyze the performance of three common configurations of energy harvesting devices as parameters are varied using numerical methods. The current literature tends to focus on analytical methods (e.g. [8]) and numerical approaches are usually reserved for analysis of random vibrations or phenomenological investigations using numerical integration. This chapter presents a numerical approach to finding the periodic orbits of the energy harvesting systems and conducting parameter studies numerically, rather than analytically. The chapter also demonstrates that uncertainty quantification is easily accomplished by numer-

ical methods. Typically, performing an uncertainty analysis analytically requires heavy symbolic computation and management of unwieldy expressions for the sensitivity coefficients. A numerical approach using finite differences for the sensitivity coefficients circumvents this difficulty entirely. By applying numerical techniques to the problem, the mathematical workload is greatly lessened for a minimal amount of additional computational effort.

4.1 Introduction

The challenge of vibrational energy harvesting in uncertain excitation environments has prompted a large movement towards nonlinear designs that broaden the usable bandwidth of energy harvesting devices [41]. Environmental excitation sources may be distributed over a wide, non-stationary spectrum [42], and as such, there has been a movement away from passive linear devices and their inherently narrow resonant peaks towards more complex devices (or systems of devices) tuned to have wider resonant peaks. One active area of research is the intentional use of nonlinearities in the harvester response to widen the usable bandwidth in an effort to better cover the excitation spectrum. A range of studies have been conducted on electromagnetic and piezoelectric systems with a variety of nonlinear restoring forces [2, 3, 4]; while each system described is quite different from the others physically, the literature has gravitated towards a common modeling framework for these nonlinear harvesters in using a modified form of the Duffing oscillator.

Recently harmonic balance was used by Stanton et al. [8] to determine the periodic orbits of the bistable electromechanical Duffing oscillator. Another recent article by Mann [7] demonstrated that harmonic balance methods can be used to conduct uncertainty analyses of the various configurations of the electromechanical Duffing oscillator. However, to the author's knowledge, only one study of the nonlinear energy harvesting system using numerical continuation has been performed to date.

Barton et al. used a continuation method to follow both the solution branches of a model and a physical experiment in [43]. Previous literature on numerical methods has focused on time-marching algorithms for general investigations of system behavior (e.g. Stanton et al. [4]) or simulation of systems undergoing stochastic forcing.

The purpose of this investigation is to demonstrate that numerical methods are a viable approach in the analysis of nonlinear energy harvesting systems. Herein we demonstrate that computation of periodic orbits can be accomplished by spectral methods, and parameter studies with sensitivity analysis can be accomplished by numerical continuation. The paper is organized as follows. First, the equations of motion for the electromechanical Duffing oscillator are reviewed. A time-periodic boundary value problem is then constructed using spectral methods to form a residual function that may be zeroed to find the periodic orbits of the system. Details regarding a linearized sensitivity analysis of the orbit are presented, as well as a method of continuing the periodic orbits and a method of computing the Floquet matrix numerically. Finally, a series of results are presented for the different configurations of the nonlinear energy harvester.

4.2 Energy harvesting device and mathematical model

While a variety of devices have been proposed for the conversion of mechanical motion to electrical power, this investigation focuses specifically on the cantilever beam designs discussed by Stanton, Erturk and others.

4.2.1 Device description and physical attributes

The purpose of this investigation is to demonstrate the merits of a structured numerical approach to analyzing energy harvesting systems in general. The analysis methods presented herein are applicable to any energy harvesting system that can

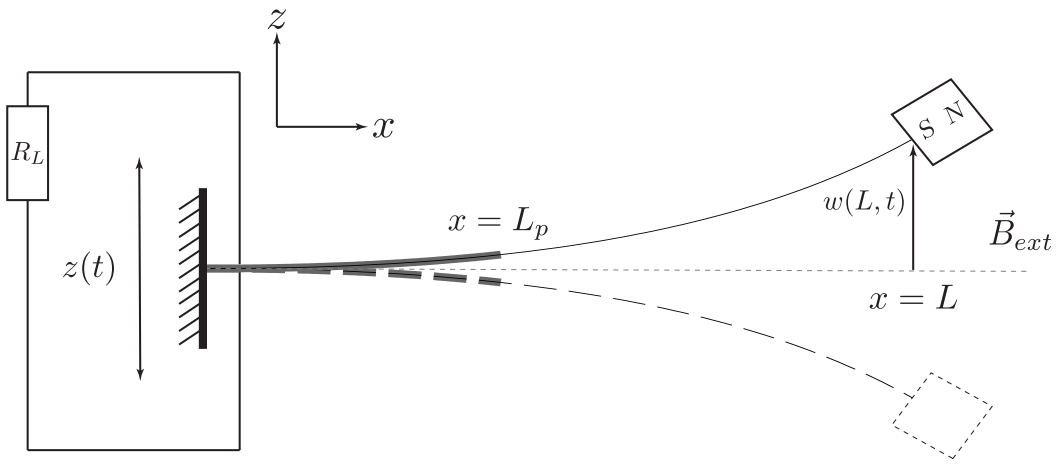

FIGURE 4.1: Schematic of an energy harvesting device consisting of a cantilevered brass beam (thin line) with a partial coating of PZT-5H (thick line). The beam has a magnetic proof mass and is in the presence of an external magnetic field $\vec{B}_{ext}$. The PZT-5H laminates are connected to an electrical load R_L.

be modeled as a system of smooth ordinary differential equations with time-periodic excitation. However, it is helpful to base the investigation on a system that could be constructed in a laboratory setting. Foundational investigations by Erturk et al. [44] and Erturk and Inman [45] focused on the performance of a partially laminated piezoelectric cantilever in a bistable configuration, and within the energy harvesting literature this design has become widely recognized and generated many subsequent investigations. Additionally, the modeling and phenomenological investigation conducted by Stanton [4] has served as a foundation for much of the mathematical modeling a device in a magnetic potential well. As such, this investigation focuses on a similar device. The model presented herein represents a brass cantilever beam, partially laminated on both sides with PZT-5H (i.e. bimorph laminates), with a magnetic end mass in an external magnetic field. The device is connected to a simple electrical circuit consisting only of a resistor and the piezolaminate layers, and the harvested voltage is measured across the resistor. A schematic of the device and circuit is shown in figure 4.1, and relevant physical parameters are shown in table

Table 4.1: Physical parameters for a nonlinear energy harvesting device.

Parameter	**Symbol**	**Value**
Brass Substrate		
Length	L_s	101.6 mm
Width	W_s	6.4 mm
Thickness	h_s	0.25 mm
Bimorph PZT-5H Laminates		
Length	L_p	25.4 mm
Width	W_p	6.4 mm
Thickness	h_p	0.27 mm
Magnetic Proof Mass		
Mass	m_t	3 g
Residual Flux Density	B_r	1.48 T

4.1. The electromechanical properties of the PZT-5H laminates are the same as in ref. [46].

An extensive survey of the electromechanical properties of this device has been conducted by Stanton, Erturk, Mann and Inman in a series of investigations [46, 47, 5]. In these investigations, it was found that the piezolaminates exhibit nonlinear coupling behavior even under moderate excitation loads [47], energy harvesting devices are intrinsically nonlinear even in the absence of nonlinear stiffness effects created by a magnetic proof mass [5], and energy harvesting devices exhibit nonlinear damping effects [46]. These investigations provide a solid foundation for a mathematical model that describes the device shown in figure 4.1.

4.2.2 Nondimensionalization and mathematical model

As detailed in the previous subsection, the model presented herein is motivated by the piezoelectric inertial generators discussed in [4, 3, 8]. Key characteristics of these devices are the presence of a relatively large end mass, which concentrates the majority of the response in the first bending mode. Thus the devices are well-approximated by either a one-mode Ritz or Galerkin approximation of the full PDE.

Following the notation and derivation presented in [8], an energy harvesting device

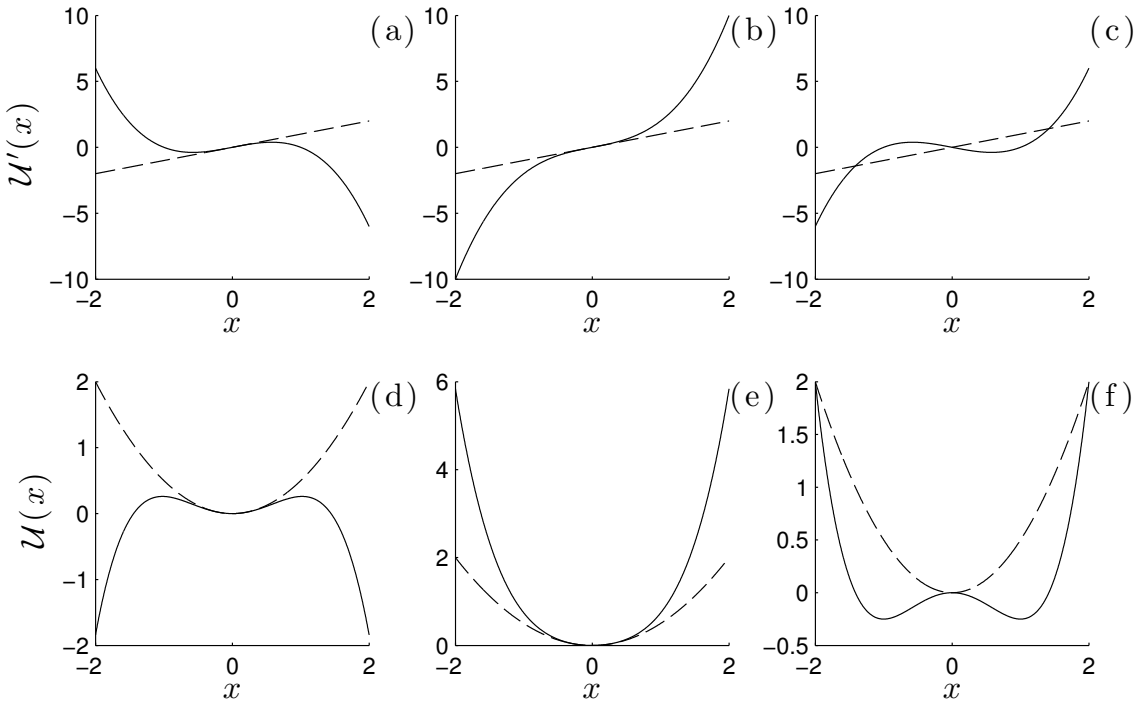

FIGURE 4.2: Sketch of the restoring forces and potentials for the (a,d) softening, (b,e) hardening and (c,f) bistable electromechanical oscillators. The top row shows the restoring forces; the bottom row shows the potential energy. The nonlinear restoring force and potential is shown by a solid line, while a linear restoring force and potential are shown by a dashed line for comparison.

with a nonlinearity that is well-represented by a cubic polynomial approximation, a capacitive circuit, and undergoing excitation at the base can be represented as

$$m\frac{\mathrm{d}}{\mathrm{d}t^2}\hat{x} + (d_a + d_b\hat{x}^2)\frac{\mathrm{d}}{\mathrm{d}t}\hat{x} + \alpha x + \beta x^3 - \Theta\hat{v} = f_z\frac{\mathrm{d}}{\mathrm{d}t^2}\hat{z} \tag{4.1a}$$

$$C\frac{\mathrm{d}}{\mathrm{d}t}\hat{v} + \frac{1}{R}\hat{v} + \Theta\frac{\mathrm{d}}{\mathrm{d}t}\hat{x} = 0. \tag{4.1b}$$

The signs of α and β determine whether the model represents a hardening (α and β both > 0), softening ($\alpha > 0$, $\beta < 0$) or bistable ($\alpha < 0$, $\beta > 0$) system. Then by substituting $\tau = t\sqrt{|\alpha|/m}$, $x = \hat{x}\sqrt{|\beta|/|\alpha|}$, $z = \hat{z}\sqrt{|\beta|/|\alpha|}$, $v = \hat{v}\sqrt{|\beta|C}/\alpha$, a

dimensionless equation

$$\ddot{x} + (\mu_a + \mu_b x^2)\dot{x} + \mathcal{U}'(x) - \theta v = \ddot{z} \tag{4.2a}$$

$$\dot{v} + \mu_c v + \theta \dot{x} = 0 \tag{4.2b}$$

can be formed, with dimensionless variables (τ, x, z, v), where $\dot{(\,)}$ indicates a derivative with respect to dimensionless time τ, and where $\mathcal{U}'(x) = -x + x^3$ for a bistable system, $\mathcal{U}'(x) = x + x^3$ for a hardening system and $\mathcal{U}'(x) = x - x^3$ for a softening system. In terms of the original physical parameters,

$$\mu_a = \frac{d_a}{\sqrt{m|\alpha|}}, \quad \mu_b = \frac{d_b}{|\beta|}\sqrt{\frac{|\alpha|}{m}}, \quad \mu_c = \frac{1}{RC}\sqrt{\frac{m}{|\alpha|}}, \quad \theta = \frac{\Theta}{\sqrt{|\alpha| C}}. \tag{4.3}$$

In this investigation, we consider harmonic forcing of the form $\frac{\mathrm{d}}{\mathrm{d}t^2}\hat{z}(t) = A\cos\Omega t$. In dimensionless form, the forcing is written

$$\ddot{z} = \Gamma \cos \eta\tau, \quad \Gamma = f_z A\sqrt{|\beta|}/|\alpha|^{3/2}, \quad \eta = \Omega\sqrt{m/|\alpha|}. \tag{4.4}$$

At this juncture, the form of the dimensionless equations matches [8]; the harmonic balance analysis from [8] will be compared with the procedure detailed in the following section to demonstrate that similar results are obtained and to highlight additional information that can be obtained from a numerical approach to the problem.

4.3 Parameter Studies with Numerical Continuation

After determining a sensible model for the energy harvesting device, the next task is often to conduct parameter studies with the model that can be used to predict the behavior of experiments. Much of the current literature has focused on analytical

methods such as harmonic balance [8, 7] and the method of multiple scales [48] to derive closed-form parameterized expressions that can be used to create response curves. These methods are excellent choices for many of the models shown in the literature so far as they yield reasonably compact symbolic expressions that can provide some intuition into the behavior of the system, even without creating graphs. However, as the modeling of nonlinear phenomena in energy harvesters becomes more detailed, these expressions become unwieldy and they lose any intuitive value.

One method, detailed herein, is numerical continuation, coupled with a numerical strategy for solving the boundary value problem for periodic orbits. To begin, the equations of motion are recast as a system of parameterized first-order differential equations

$$\dot{\boldsymbol{x}} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_k \end{pmatrix} = f(\boldsymbol{x}, t; p) = \begin{pmatrix} f_1(\boldsymbol{x}, t; p) \\ f_2(\boldsymbol{x}, t; p) \\ \vdots \\ f_K(\boldsymbol{x}, t; p) \end{pmatrix}, \tag{4.5}$$

where $\boldsymbol{x} = (x_1, x_2, \ldots, x_K)^\top \in \mathbb{R}^K$ are the states, $f_k(\boldsymbol{x}, t; p) : \mathbb{R}^{K+1} \to \mathbb{R}$ are the K non-autonomous state equations and p is the set of physical parameters. For example, the bistable system would be represented as:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{pmatrix} = \begin{pmatrix} x_2 \\ -(\mu_a + \mu_b x^2)\dot{x} + x - x^3 + \theta x_3 + \Gamma \cos \Omega t \\ -\mu_c x_3 - \theta x_2 \end{pmatrix} \tag{4.6}$$

where x_1 is the dimensionless displacement of the harvester, x_2 is the dimensionless velocity and x_3 is the dimensionless voltage, with the set of parameters $p = \{\mu_a, \mu_b, \mu_c, \theta, \Gamma, \Omega\}$. Given this first order system, closed orbits are found by solving

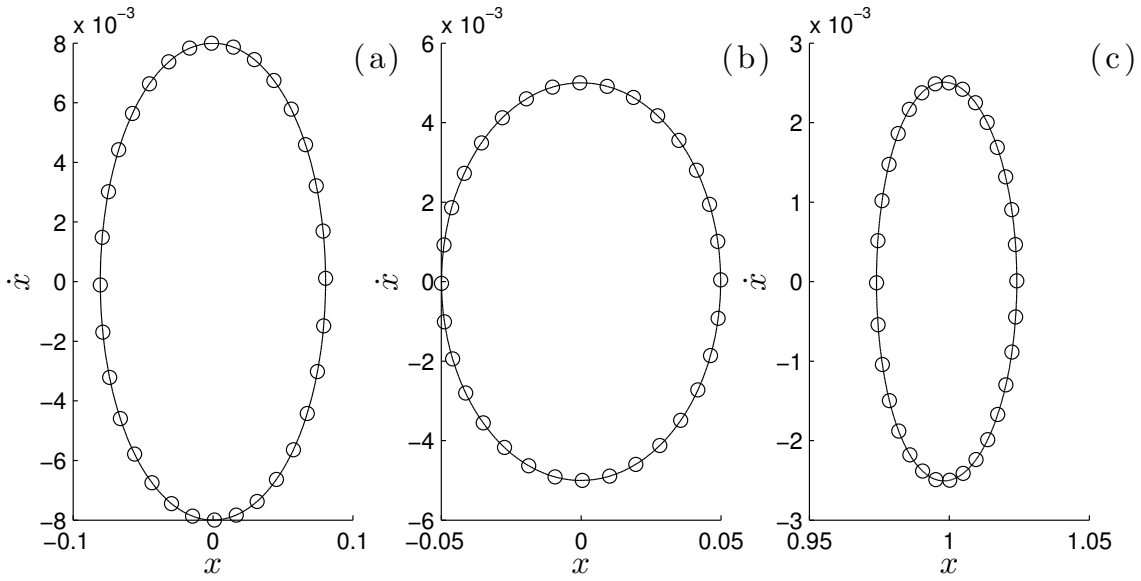

FIGURE 4.3: Comparison of numerical integration (-) and collocation (○) solutions for the softening (a), hardening (b) and bistable (c) oscillators for a stable orbit. Parameter values used in all cases are $\mu_a = 0.08$, $\mu_b = 0.01$, $\mu_c = 0.01$, $\theta = 0.1$, $\Gamma = 0.08$ and $\eta = 0.1$. The agreement is nearly exact for stable solutions; unstable solutions diverge gradually. The equal spacing of the collocation points in time is clearly visible here.

the periodic boundary value problem

$$\begin{pmatrix} \dot{\boldsymbol{x}} \\ 0 \end{pmatrix} = \begin{pmatrix} Tf(\boldsymbol{x}, \tau; p) \\ \boldsymbol{x}(1) - \boldsymbol{x}(0) \end{pmatrix} \tag{4.7}$$

where premultiplying $f(\boldsymbol{x}; p)$ by the period of the orbit T rescales $t \in [0, T] \to \tau \in [0, 1]$ and the overdot now signifies a derivative with respect to τ. The boundary condition $\boldsymbol{x}(1) = \boldsymbol{x}(0)$ enforces periodicity. This time-periodic BVP (TP-BVP) may now be discretized and solved numerically.

4.3.1 Solution of the TP-BVP by psuedospectral collocation

To formulate the numerical boundary value problem, the orbit is discretized into a set of n collocation points $\boldsymbol{\tau} = (\tau_0, \tau_1, \ldots, \tau_n)^\top \in \mathbb{R}^n$, a set of $n \times k$ sample states $y = (\boldsymbol{x}^\top(\tau_0), \boldsymbol{x}^\top(\tau_1), \ldots, \boldsymbol{x}^\top(\tau_n))^\top \in \mathbb{R}^{n \times k}$ and a differentiation matrix $\boldsymbol{D} \in \mathbb{R}^{n \times n}$ is

created such that

$$\begin{pmatrix} \dot{\boldsymbol{x}}^\top(\boldsymbol{\tau_0}) \\ \dot{\boldsymbol{x}}^\top(\boldsymbol{\tau_1}) \\ \vdots \\ \dot{\boldsymbol{x}}^\top(\boldsymbol{\tau_n}) \end{pmatrix} \approx \boldsymbol{D} \begin{pmatrix} \boldsymbol{x}^\top(\boldsymbol{\tau_0}) \\ \boldsymbol{x}^\top(\boldsymbol{\tau_1}) \\ \vdots \\ \boldsymbol{x}^\top(\boldsymbol{\tau_n}) \end{pmatrix} + \mathcal{O}(\varepsilon) \tag{4.8}$$

or more compactly $\dot{y} = \boldsymbol{D}y + \mathcal{O}(\varepsilon)$. Rows of y, $y_i = (x_1(\tau_i), \ldots, x_K(\tau_i))$, correspond to samples at the collocation points τ_i while columns correspond to states x_k; $\boldsymbol{D}$ is a square matrix whose rows correspond to the derivative approximation at τ_i and whose columns correspond to the influence of neighboring samples $y(\tau_i)$ on the approximation of the derivative. The matrix $\boldsymbol{D}$ can be constructed in a number of ways; a Fourier spectral differentiation matrix is chosen for this investigation because it enforces the periodic boundary conditions explicitly, but other matrices such as a finite difference, finite element, or spectral element matrix could be used, as well as a different set of basis functions for determining the collocation points such as Chebyshev or Legendre polynomials. See [49] or [50] for numerous examples as well as details for generating $\boldsymbol{D}$.

Assuming equation (4.8) is valid, $\boldsymbol{D}y = Tf(y, \boldsymbol{\tau}; p)$ is a valid approximation of equation (4.7). We can then write a residual for the boundary value problem as

$$R(y; p) = \boldsymbol{D}y - Tf(y, \boldsymbol{\tau}; p)^\top, \tag{4.9}$$

where the substitution $\dot{y} \approx \boldsymbol{D}y$ has been made[1]. In a collocation method, zeros of the residual function correspond directly to discretized solutions (periodic orbits) of equation (4.7); other methods (Galerkin, etc.) have a slightly different interpretation (c.f. [50, 17]). An example solution of equation (4.9) is shown in Figure 4.3.

[1] As a convention throughout, any arguments that follow a semicolon are assumed to be fixed during the solution of the TP-BVP. In the energy harvesting context, examples would be damping, forcing frequencies or amplitudes, or coupling parameters.

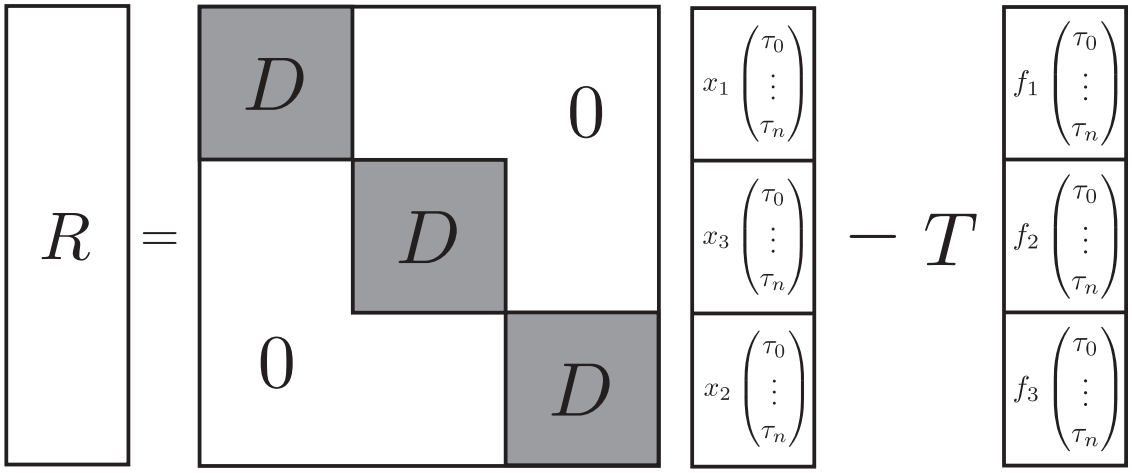

FIGURE 4.4: Schematic of the vertically concatenated residual for a system with three states.

To enforce the boundary condition $y(0) = y(1)$, the last row of D, which corresponds to $\tau = 1$ is typically replaced with $(1, 0 \ldots, 0, -1)$ to enforce periodicity (typically referred to as "boundary bordering the matrix"); however, when using a Fourier spectral differentiation matrix periodicity is implicitly enforced and the last row and column (corresponding to $\tau = 1$) are removed to avoid problems of collinearity. See [49] for an overview of including boundary conditions and [50] for a more detailed discussion of boundary conditions.

Let

$$R_{\tau_i}(x_k; p) = \boldsymbol{D}_{\tau_i} \begin{pmatrix} x_k(\tau_0) \\ x_k(\tau_1) \\ \vdots \\ x_k(\tau_n) \end{pmatrix} - T f_k((x_1(\tau_i), \ldots, x_k(\tau_i))^\top, \tau_i; p) \tag{4.10}$$

be the residual at the i^{th} collocation point τ_i for the k^{th} state x_k with D_{τ_i} the i^{th} row of the differentiation matrix D. Hence $R_{\tau_i}(x_k; p)$ is scalar and corresponds to the error at collocation point $x_{i,k}$ given all other collocation points of the same state

and all states at the same time. Then the total residual matrix has the form

$$R(y;p) = \begin{pmatrix} R_{\tau_1}(x_1;p) & \dots & R_{\tau_1}(x_K;p) \\ \vdots & \ddots & \vdots \\ R_{\tau_n}(x_1;p) & \dots & R_{\tau_n}(x_K;p) \end{pmatrix} \tag{4.11}$$

with $R(y;p) \in \mathbb{R}^{n\times K}$. This form is compact, but computationally inefficient due to the matrix-matrix multiplication of D and y. It is advantageous to vertically concatenate the columns of y and $f(y, \boldsymbol{\tau}; p)$ and instead write the residual

$$R(y;p) = (I_k \otimes \boldsymbol{D})\,\mathrm{vect}\{y\} - T\,\mathrm{vect}\{f(y,\tau;p)^\top\} \tag{4.12}$$

where $\otimes$ is the Kronecker product and I_K is the $K \times K$ identity matrix (3×3 in the case of the 3-state energy harvester model). Now we have $R(y;p) \in \mathbb{R}^{n\cdot K}$. A schematic of the concatenated residual structure is shown in figure 4.4. This vector residual is computationally preferable to the matrix residual above as a modified Newton method will be used to solve $R(y;p) = 0$ to find periodic orbits.

The final remaining question is how to determine the appropriate number of collocation points: for this investigation, a solution generated by collocation was compared to a solution generated by numerical integration, as seen in Figure 4.3, for varying numbers of collocation points. As can be seen in Figure 4.5, the error falls off exponentially as the number of points used increases. For this investigation, 32 collocation points[2] were used to balance error against computational expense. This formulation of the TP-BVP is used to compute periodic orbits, conduct uncertainty analyses and follow the orbits as parameters are varied using numerical continuation.

[2] Note that selecting a number of collocation points $n = 2^m$, where m is a whole number, creates a problem that is easily parallelizable across multiple CPU cores.

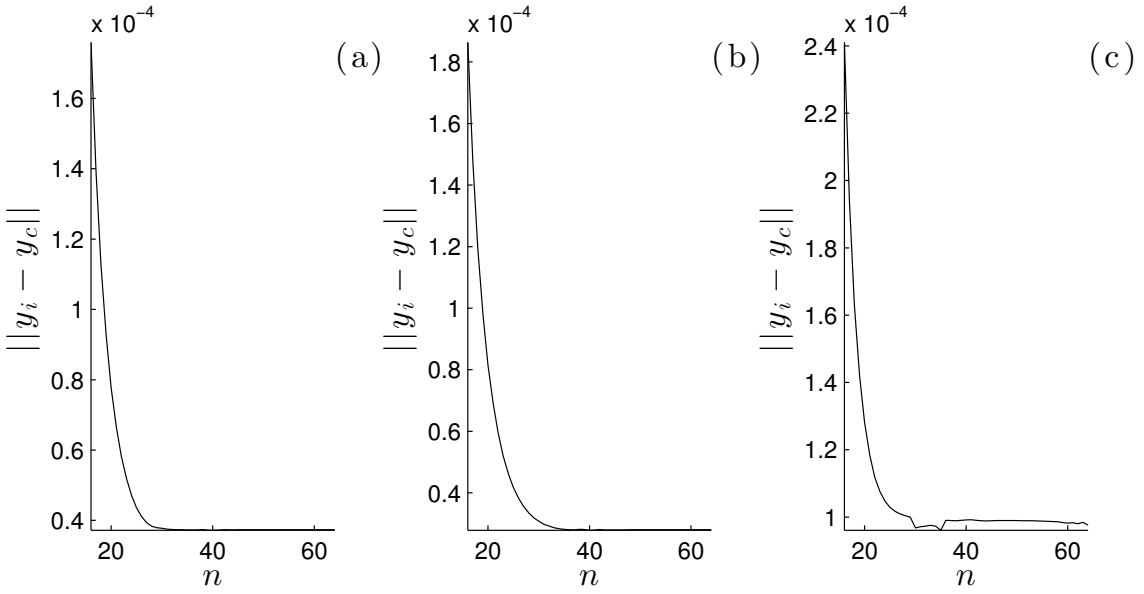

FIGURE 4.5: Convergence of error $||y_i - y_c||$ in collocation solution y_c relative to a numerically integrated trajectory y_i for the softening (a), hardening (b) and bistable (c) cases where n is the number of collocation points.

4.3.2 Linearized uncertainty analysis of a periodic orbit

As discussed in [7], determining the sensitivity of orbits to variance in system parameters is an important step in determining real-world limits on performance. Here, a linearized uncertainty analysis is performed as discussed in [51] and used in [7] in conjunction with harmonic balance to study the energy harvester equations. Note that a more detailed method such as generalized polynomial chaos could be used, but for proof of concept, linearized uncertainty analysis is satisfactory. Those interested in more powerful methods such as generalized polynomial chaos should consult ref. [52]. Lacking an analytical solution, the uncertainty analysis must be performed numerically.

To begin, assume that the solution y which solves $R(y;p) = 0$ is a function of the

P parameters p, that is

$$y(p) = y(p_1, p_2, \ldots, p_P). \tag{4.13}$$

Then, assume there is an uncorrelated, normally distributed perturbation $(\epsilon_1, \epsilon_2, \ldots, \epsilon_n)$ with each ϵ_i representing the uncertainty in each p_i; the assumption of normality is used to make the analysis tractable [51]. The solution with stochastic parameters can be written as $y_s = y(p_1 + \epsilon_1, p_2 + \epsilon_2, \ldots, p_P + \epsilon_P)$. The mean and variance of S are given by

$$\mu_S = \mathbb{E}\left[y_s\right] = y(p_1, p_2, \ldots, p_P) \tag{4.14}$$

$$\sigma_S^2 = \mathbb{E}[y_s^2] - \mathbb{E}[y_s]^2 = \sum_{i=1}^{P} \left(\frac{\partial y}{\partial p_i}\right)^2 \mathbb{E}[\epsilon_i^2], \tag{4.15}$$

which can be found by Taylor expanding and canceling terms (see [51]). Note that $\sigma_{p_i}^2 = \mathbb{E}[\epsilon_i^2]$ is simply the variance of ϵ_i. Now, define $U_{p_i} = n\sigma_{p_i}$ as the uncertainty in parameter p_i, and

$$\theta_i^2 = \left(\frac{\partial y}{\partial p_i}\right)^2 \tag{4.16}$$

as the sensitivity coefficient corresponding to p_i. If n is chosen to be an integer it will correspond to a confidence interval at n standard deviations. Then the total uncertainty in the solution U_y is given by

$$U_y^2 = \sum_{i=1}^{P} \theta_i^2 U_{p_i}^2. \tag{4.17}$$

Lacking an analytical expression for $S(p)$, we determine the coefficients θ_i, with

a simple forward difference, that is

$$\theta_i = \frac{y(p+\delta) - y(p)}{\delta}, \tag{4.18}$$

where δ is a small perturbation applied to p_i. In this investigation we use $\delta = 0.005 \cdot p_i$ (i.e. a 0.5% perturbation in p_i) to balance accuracy against roundoff errors. The solutions $y(p)$ and $y(p+\delta)$ are found by solving $y(p) = \{y \mid R(y;p) = 0\}$ and $y(p+\delta) = \{y \mid R(y;p+\delta) = 0\}$.

4.3.3 Numerical continuation of periodic orbits

In the context of energy harvesting, an investigator is most often interested in obtaining the response of a system as forcing frequencies and amplitudes are varied, as well as determining the influence of tuning parameters on the behavior of the system. Thus, computing the response trends of the TP-BVP to variations in a parameter is fundamental to any useful energy harvesting investigation. To do so, begin by finding a reference solution $y_0 = \{y \mid R(y, \alpha_0; \hat{p}) = 0\}$ with a fixed subset of parameters $\hat{p} \subset p$ and one bifurcation parameter α chosen from the full set of parameters p. For example, one might choose α to be the forcing frequency; $\hat{p}$ would be the remainder of the parameters such as the forcing amplitude, electromechanical coupling and dissipative parameters. To avoid repetitive notation, let $R(y, \alpha; \hat{p})$ be $R(y, \alpha)$.

Given y_0, one then seeks to follow the branch of solutions on which y_0 lies over a certain range of α. For a TP-BVP strategy with n collocation points, numerical continuation can be used to find a discretized representation of the n-dimensional manifold $\mathcal{M}$ embedded in the $n+1$ dimensional (y, α) space of collocation variables

y and the continuation parameter α. $\mathcal{M}$ is defined as

$$\mathcal{M} = \{(y, \alpha) \mid R(y, \alpha) = 0, \alpha \in [\alpha_{\min}, \alpha_{\max}], \exists \text{ a path from } y_0 \text{ to } y\}, \tag{4.19}$$

i.e. the zero-set of $R(y, \alpha)$ in the parameter range $[\alpha_{\min}, \alpha_{\max}]$ that is connected to y_0 [53]. Note that in general numerical continuation algorithms will not find disconnected branches without user intervention. In the energy harvesting context, an example of disconnected branches would be the inter- and intra-well solution curves of the bistable oscillator.

It is important to note that the manifold $\mathcal{M}$ is embedded in the $n+1$ dimensional space that includes both collocation variables y and the continuation parameter α: $\mathcal{M}$ is an implicit function of both. A naive strategy of marching incrementally from $\alpha_{\min}$ to $\alpha_{\max}$ would fail at points where the Jacobian $R_y(y, \alpha)$ becomes singular. A physical example of this would be at the points in a nonlinear frequency response function when the slope becomes vertical. Hence an additional degree of freedom is included by allowing α to become a free parameter and adding a constraint equation to keep the system square. Let $\mathcal{M}_d = \{(y_0, \alpha_0), \ldots, (y_j, \alpha_j), \ldots\}$ be the discretized manifold. Then by repeatedly solving

$$\begin{pmatrix} 0 \\ ds \end{pmatrix} = \begin{pmatrix} R(y, \alpha) \\ ||(\Delta y, \Delta \alpha)||_2 \end{pmatrix} \equiv R^+(y, \alpha) \tag{4.20}$$

a collection of (y_j, α_j) can be found that represent the manifold. In equation (4.20) $\Delta y = y_j - y_{j-1}$ and $\Delta \alpha = \alpha_j - \alpha_{j-1}$ are the vectors from the previous solution for y and α and the final additional equation $||(\Delta y, \Delta \alpha)||_2 = ds$ enforces a step length of ds. Including this final equation ensures that the Jacobian $R^+_y(y, \alpha)$ does not become singular at turning points and keeps the system square with the addition of the extra free variable α. This choice of a constraint equation is by no means unique; other

choices are detailed in [24, 23, 25]. A secant predictor

$$\hat{y} = y_{j-1} + (y_{j-1} - y_{j-2}) \tag{4.21}$$

$$\hat{\alpha} = \alpha_{j-1} + (\alpha_{j-1} - \alpha_{j-2}) \tag{4.22}$$

is used as the initial guess for each new solution y_j and a trust-region dogleg algorithm is used to find the zeros of $R^+(y, \alpha)$.

4.3.4 Numerical bifurcation and stability analysis

Finally, a method is required to detect any bifurcations that might occur in the system. Again, lacking analytical expressions for the orbits, numerical methods must be used. To find a bifurcation, a test function $\phi(y, \alpha)$ is constructed such that $\phi(y, \alpha) = 0$ at a bifurcation [23, 25, 24]. This test function is constantly monitored during the continuation process, creating a series of test function values $\phi_j = \phi(y_j, \alpha_j)$ corresponding to the pairs (y_j, α_j) of $\mathcal{M}_d$. Then the condition

$$\phi_{j-1}\phi_j < 0 \tag{4.23}$$

indicates that the test function changed signs between steps j and $j-1$. At this point, the continuation is halted and a new step from $j-1$ is made with the added condition of zeroing the test function, that is:

$$\begin{pmatrix} 0 \\ 0 \end{pmatrix} = \begin{pmatrix} R^+(y, \alpha) \\ \phi(y, \alpha). \end{pmatrix} \tag{4.24}$$

This finds the bifurcation point directly. Either a Newton-type method or a secant-type method can be used to find the bifurcation point; as the Jacobian of eq. (4.24) is sometimes expensive to compute, secant-type methods are sometimes preferable despite their slower convergence than Newton methods. In this investigation we consider three types of bifurcations: the turning point bifurcation, given by

$$\phi_{TP} = \det(R_y(y, \alpha)), \tag{4.25}$$

where $R_y(y, \alpha)$ is the Jacobian of the collocation residual; the period-doubling bifurcation, given by

$$\phi_{PD} = \det(\Phi - I_K), \tag{4.26}$$

where Φ is the monodromy matrix; and the symmetry-breaking bifurcation, given by

$$\phi_{SB} = \det(\Phi + I_K). \tag{4.27}$$

All of these test functions are can be found in either [24, 23] or [25]. Calculation of $R_y(y, \alpha)$ is straightforward and in practice one simply uses all but the last row and column of $R_y^+(y, \alpha)$, as $R_y^+(y, \alpha)$ is calculated in the process of solving equation (4.20). Determining the monodromy matrix Φ is less straightforward. To find Φ, the system

$$\begin{pmatrix} \dot{\boldsymbol{x}} \\ \dot{\Phi} \end{pmatrix} = \begin{pmatrix} f(\boldsymbol{x}, t; \langle \alpha, \hat{p} \rangle) \\ f_{\boldsymbol{x}}(\boldsymbol{x}, t; \langle \alpha, \hat{p} \rangle)\Phi \end{pmatrix} \tag{4.28}$$

is numerically integrated with initial conditions $\boldsymbol{x}(0)$ taken from the current orbit, computed in equation (4.20), and $\Phi(0) = I_k$ (the $k \times k$ identity matrix). In practice, the matrix ODE for Φ is reshaped into a vector to ensure compatibility with numerical routines. Solving this coupled $(k + k^2)$-state ODE on $t = [0, 1]$ yields $\Phi(1) \equiv \Phi$, the monodromy matrix. Monitoring the eigenvalues of Φ allows us to detect bifurcations: when an eigenvalue μ_i of Φ satisfies $|\mu_i| = 1$, a bifurcation occurs. This investigation considers only two (period-doubling and symmetry-breaking) of the three types of bifurcations related to the monodromy matrix in this investigation; the third, Neimarck-Sacker bifurcations, were not detected in the systems studied. This is typically the most expensive computation during each step as equation (4.20) usually converges within a few iterations if a reasonable step size is maintained.

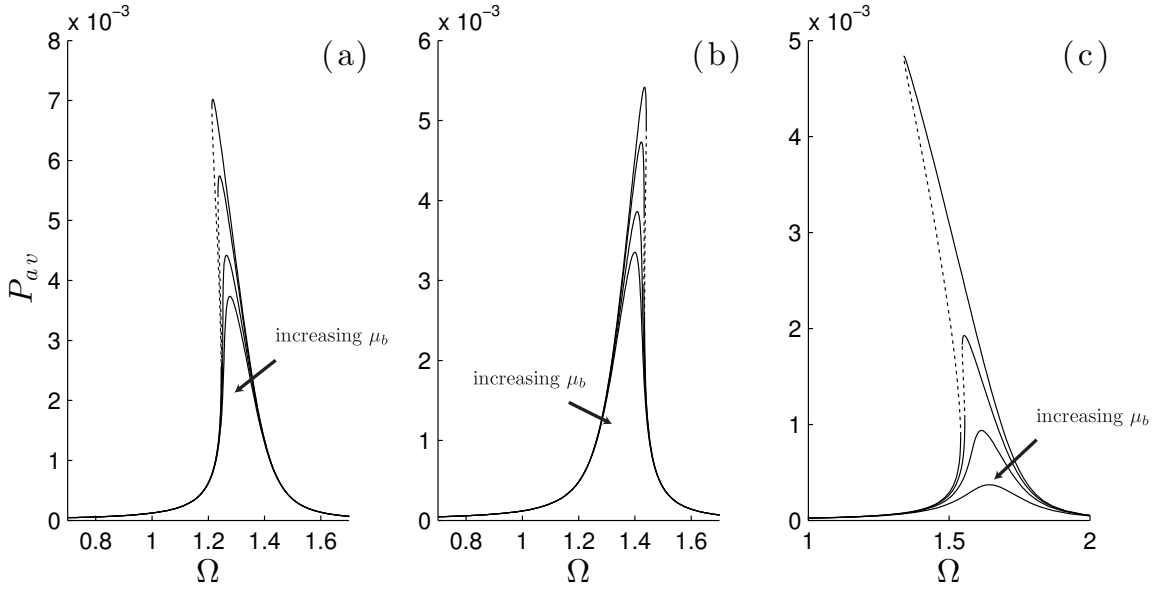

FIGURE 4.6: Response trends in average power (P_{av}) per cycle illustrating the suppressing effects of increasing μ_b for the softening (a), hardening (b) and bistable (c) configurations. Other parameter values are $\mu_a = 0.08$, $\mu_c = 0.01$, $\Gamma = 0.07$ and $\theta = 0.9$. Stable solutions exist on the solid lines; dotted lines indicate unstable solutions.

4.4 Results

In the previous section, methods for analysis of a mathematical model of a harmonically-excited energy harvester were detailed, including: obtaining a periodic orbit through the solution of a periodic boundary value problem via psuedospectral collocation; determining the sensitivity of the orbit to uncertainties in the model parameters; following the solution curve of the model as parameters are varied; determining the stability of orbits and identifying bifurcations on the solution curves. While the process detailed above differs in approach from much of the existing literature, the fundamental goal is the same: determining performance trends in various excitation environments and determining the effects of tunable parameters on the device's performance.

In this section, we implement the procedures discussed in the previous section to

conduct three series of investigations. First, we observe the influence of the damping parameters and the electromechanical load on the overall response trends of each system; these results are compared to [8] to emphasize the ability of the collocation method to obtain similar results to harmonic balance. Second, we illustrate the effects of uncertainties in the parameters; these results are compared to [7] to show that calculating the sensitivities numerically yields similar results to the analytical methods used in [7]. Finally, we investigate the bifurcations in each system and observe behaviors that were not found by harmonic balance.

4.4.1 Response trends and energy harvesting implications

In [8], response trends were shown as the nonlinear damping μ_b and electrical load μ_c parameters were varied, highlighting the effects of these parameters on the behavior of the bistable system. Here we consider all three configurations and illustrate the response trends as μ_b and μ_c are varied.

In figure 4.6, the suppressing effect of the nonlinear damping parameter μ_b can be clearly seen. With a sufficiently high nonlinear damping parameter, the cusp bifurcation causing a loss of stability and multiple coexisting solutions disappears in all three systems. This effect is consistent for both high and low amplitude excitations and matches with the results found in [8].

Varying the electrical load μ_c has a more noticeable effect on the resonant peak of the system. Increasing the electrical load decreases the amplitude of the response but increases the average power (P_{av}) up to a certain limit. Additionally, the resonant peak of the device shifts in frequency as μ_c is varied, again in all three configurations. This can be seen in figure 4.7. This is an important consideration in the design of a device because in addition to being an easily tunable parameter, the circuit resistance alters the resonant peak and bandwidth more drastically than any of the other parameters.

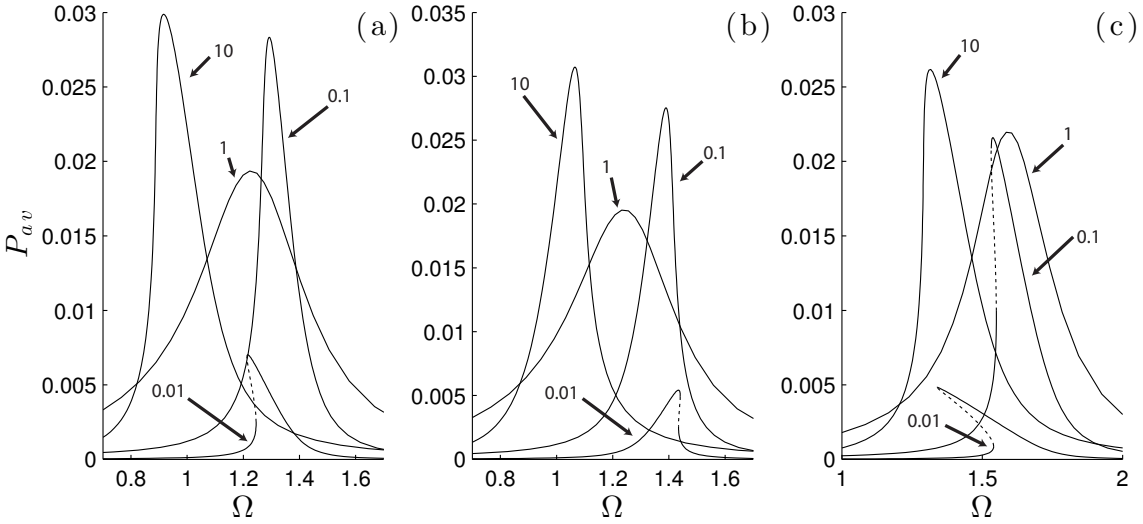

FIGURE 4.7: Response trends in average power (P_{av}) per cycle for electrical loads $\mu_c = \{0.01, 0.1, 1, 10\}$ for the softening (a), hardening (b) and bistable (c) configurations. Other parameter values are $\mu_a = 0.08$, $\mu_b = 0.01$, $\Gamma = 0.07$ and $\theta = 0.9$. Stable solutions exist on the solid lines; dotted lines indicate unstable solutions.

Figure 4.8 illustrates the nature of the system response with respect to the electrical load. For each system, there is a clear peak in power output. This peak shifts as the frequency or amplitude of excitation shifts, indicating that determining an optimal circuit load is a nontrivial optimization problem.

4.4.2 Bifurcation behavior

In the previous subsection, response trends were shown that demonstrated the effects of nonlinear damping and the electrical load on the behavior of the system. In this section, results are presented that focus on the bifurcations that occur in the system as the forcing amplitude is varied.

It is well known that Duffing-type oscillators exhibit period-doubling and symmetry breaking bifurcations as forcing amplitude is increased. Similar behavior is observed in all three configurations of the energy harvesting device. Figure 4.9 shows trends as the forcing amplitude is increased near the resonant peak of each device. Notably, each device exhibits coexisting solutions, as has been reported in several

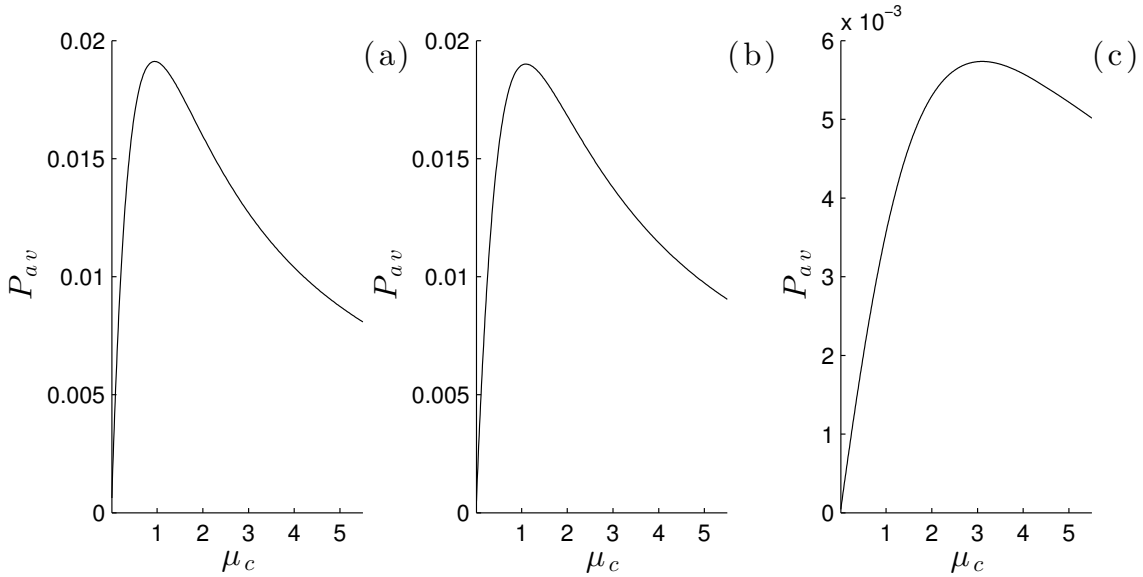

FIGURE 4.8: Response trends of average power per cycle (P_{av}) as the electrical load μ_c is varied. In each case, there is a clear peak, highlighting the presence of an optimal impedance value for each harvester design. The softening case is shown in (a), hardening in (b) and bistable in (c). Parameters shown are $\mu_a = 0.08$, $\mu_b = 0.01$, $\theta = 0.9$, $\Gamma = 0.07$ and $\Omega = 1.2$.

prior investigations. In this investigation, we identified several period-doubling and symmetry-breaking bifurcations and found evidence of physical behavior that has not yet been thoroughly discussed in the literature to the best of our knowledge.

For each configuration of the device, a small-amplitude solution persists through moderate values of Γ, the forcing amplitude. The solution coexists with a higher-amplitude solution that would typically be unstable in a purely mechanical Duffing-type oscillator but is stable in the electromechanical version shown here, as confirmed by numerical simulation and a test on the Floquet multipliers of the system. As the Γ is increased further, the low-amplitude solution ceases to exist altogether and all three configurations go through a cascade of period-doubling bifurcations intermixed with symmetry breaking bifurcations.

The presence of a symmetry-breaking bifurcation indicates the formation of a solution with a winding number $\neq 1$: that is to say, the solution folds over itself

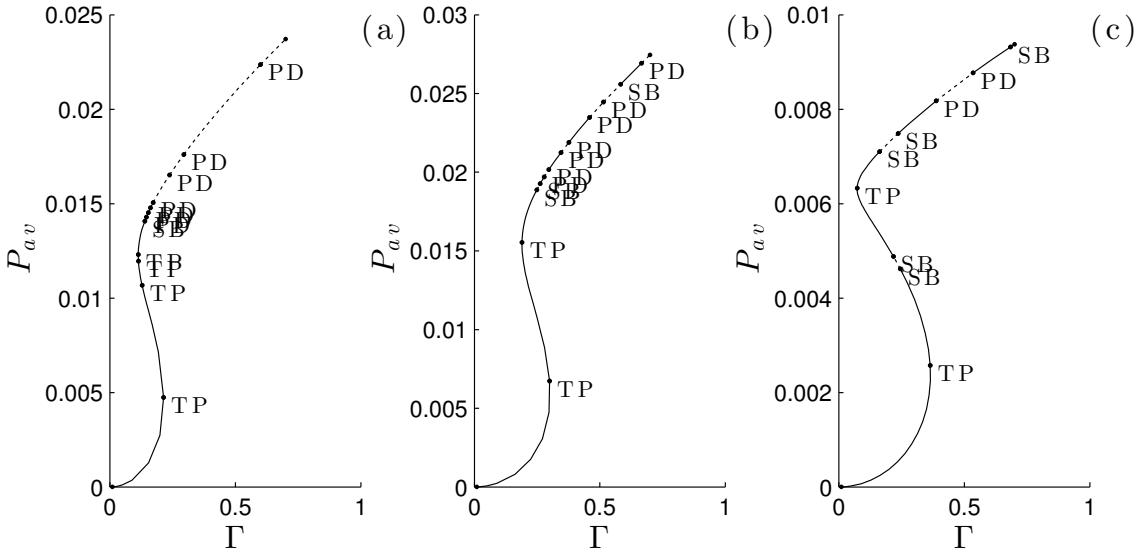

FIGURE 4.9: Bifurcations observed as the forcing amplitude is increased near the resonance of each device. TP indicates a turning point or cusp-type bifurcation, SB indicates a symmetry breaking bifurcation, and PD indicates a period doubling bifurcation. PL $\in \{\Gamma_{min}, \Gamma_{max}\}$ represents a parameter limit that terminated the continuation algorithm.

when viewed in a two-dimensional phase plane and is no longer symmetric about the displacement axis, and the solution is no longer homeomorphic to the circle. Period-doubling solutions often appear visually similar to symmetry breaking solutions, but indicate the presence of a symmetric orbit of doubled period. In each system, the presence of period doubling bifurcations causes a different phenomenon to occur.

In the softening system, period-doubling bifurcations occur in rapid succession as the forcing intensity nears $\Gamma \approx 0.2$, the system undergoes a well escape and the amplitude of the response goes to infinity. This is of course a non-physical response and represents an inherent shortcoming in the modeling of the restoring force. Had the restoring force been truncated at fifth order rather than third order, the potential well for the softening configuration (as seen in figure 4.2) would resemble a flattened parabola rather than an inverted bistable potential well, thus avoiding the nonphysical numerical/modeling problem of well escapes altogether. A clearer

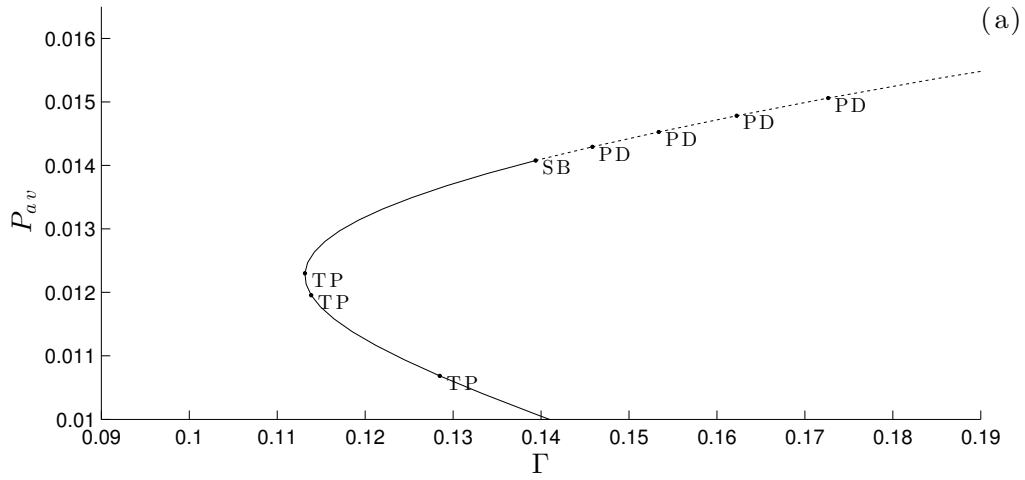

FIGURE 4.10: Period doubling cascade observed in the softening device approaching the well escape. A drawback of the truncating the stiffness approximation at third order is the manifestation of non-physical well escape behavior in the softening system. TP indicates a turning point or cusp-type bifurcation, SB indicates a symmetry breaking bifurcation, and PD indicates a period doubling bifurcation.

view of the bifurcation sequence can be seen in figure 4.10.

In the hardening system, the succession of period-doubling bifurcations are less meaningful as the system alternately loses and regains a stable period-1 orbit as the forcing intensity is increased. The hardening configuration is by far the most robust in terms of maintaining stable period-1 orbits even under high forcing amplitudes and near resonant conditions. This is potentially advantageous as it eliminates the need for complicated circuit designs to efficiently harvest power from the device.

In the bistable system, the period doubling bifurcations again correspond to a well escape, but unlike the softening system the well escape in the bistable system has a physical meaning corresponding to the transition from single-well to cross-well solutions. The intrawell solution (represented by the branch shown in figure 4.9c) is intermittently stable even through very high forcing amplitudes. This highlights an important challenge in the design of a bistable harvester that many others have already noted in the literature - the low-energy intrawell solution persists even at high

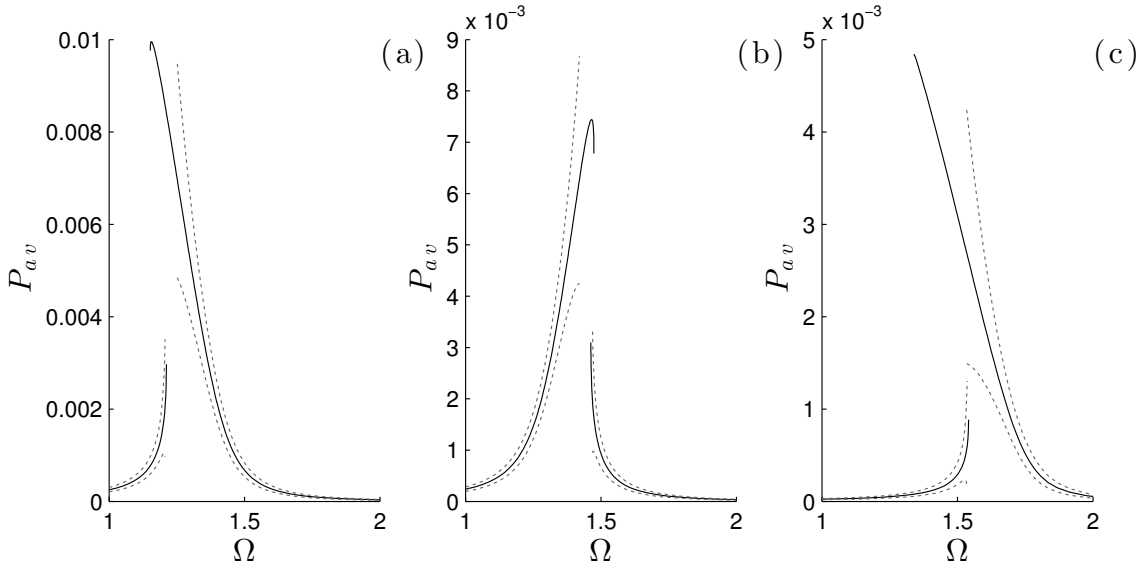

FIGURE 4.11: Uncertainty in the frequency response of the softening (a), hardening (b) and bistable (c) configurations. The mean response is shown by a black line; the 95% confidence interval is shown by gray dotted lines. The uncertainty is not shown when it approaches numerical infinity; additionally unstable branches of the response are not shown.

amplitudes and maintaining a high-energy interwell solution would usually require intervention through a control scheme.

4.4.3 Uncertainty quantification

In a physical system, parameters can never be identified exactly. There is always some measure of uncertainty in the measured parameters, and if the measurement uncertainty is small, the linearized sensitivity analysis detailed above is an excellent way of determining the influence of these uncertainties on the performance of the device. In figure 4.11, the general trend of response uncertainty is clearly visible. There is little variance in the response away from the resonant peak, but near the resonant peak, all three devices show considerable variance in the response magnitude. Specifically, the bistable device has the greatest relative uncertainty in response near resonance, followed by the softening device. The hardening device has a fairly

robust response with low variance near resonance. Given the uncertain nature of energy harvesting environments, computing the response uncertainties provides useful information to the investigator, and by using the method described in the previous section, the uncertainties are trivial to calculate while in the process of calculating the mean response of the system.

4.5 Conclusions

This investigation has demonstrated the merits of a numerical approach to characterizing the response of energy harvesting systems to harmonic forcing. While many prior investigations have utilized numerical integration to characterize system behavior, the approach detailed herein allows the systems to be studied in greater detail, including finding unstable orbits, finding bifurcations of orbits, and finding the sensitivity of orbits to uncertainties in physical parameters. The methods presented allow for the influence of all system parameters to be investigated with ease. Additionally, the methods presented in this investigation are not computationally intensive compared to a brute-force search of the parameter space with numerical integration; it takes tenths of a second to find a point on a solution curve, check its stability, conduct a sensitivity analysis and monitor for bifurcations, allowing entire solution curves to be built in less than a minute in most cases. The method also circumvents difficulties that might arise in symbolic computation while finding approximate analytical solutions, instead allowing the investigator to conduct parameter studies from the original system of ODEs directly. However, there is still great value in obtaining symbolic solutions for the intuition they provide into the behavior of a system.

Future numerical continuation work on energy harvesting systems should focus on devices which either do not admit approximate analytical solutions, or are beyond the level of complexity at which a reasonably compact analytical solution can be obtained. This will allow for more thorough investigations of the wide range of

nonlinear behaviors present in these devices. Additionally, stochastic forcing and higher-fidelity uncertainty quantification should be considered using the methods presented in [52]. Branching behavior at bifurcation points should be considered and solution curves emanating from the bifurcation points should be considered; for a theoretical overview see [23]. To aid in the visualization of response trends against multiple parameters, the multiparameter continuation method detailed in [53] could be used. Finally, continuation techniques should be employed in a laboratory setting as demonstrated in [43] and combined with a real-time controller to follow the solution branches of physical devices for comparison with mathematical models.

5

Phenomenological Investigation of a Tilt-Excited Pendulum

As discussed in chapter 1, the construction of an effective energy harvester in an ocean environment almost certainly necessitates the use of a nonlinear device. While linear devices benefit from a consistent, predictable response, even in random excitation environments, they are fundamentally limited by a narrow bandwidth from which energy can be harvested. Nonlinear devices lack the predictability of linear devices, but they often exhibit highly energetic behaviors over a wider range of the frequency spectrum than their linear counterparts.

One such nonlinear device is a proof mass affixed to a generator via a lever arm on a rotating bearing. As the buoy or vessel heaves and pitches in the ocean waves, the changing tilt of the apparatus relative to the vertical axis induces motion by altering the gravitational potential energy of the proof mass. Energy is then harvested when the lever arm rotates and turns the generator, converting the kinetic energy of the proof mass into electrical energy in the harvesting circuitry.

A significant advantage of this design is its proven history as one of the oldest and most persistent kinetic energy harvesting technologies. The principle of a rotating

proof mass driven by tilt-induced changes in gravitational potential is a well-tested technology that has been employed in self-winding, or automatic, watches since the eighteenth century. In the 1770's, Abraham-Louis Perrelet is said to have invented the first self-winding watch mechanism for a pocketwatch, utilizing a pendular device excited by the motion of the watches user to incrementally wind the mainspring [9]. His mechanism was later improved by Abraham-Louis Breguet, but the limited motion experienced by the watches in their owner's pockets limited the commercial success of the devices [9].

The self-winding wristwatch was not invented until the twentieth century, when John Harwood developed the "bumper" watch. Patented in 1923, the design employed a pendular mechanism with a limited range of motion that encouraged librating behaviors to gradually wind the mainspring [9]. Due to the increased motion experienced at the wrist, the self-winding mechanism functioned well and the watch was a commercial success until the Great Depression. The first design to incorporate a self-winding mechanism capable of unrestricted rotation was the Rolex Oyster Perpetual, which employs a semicircular proof mass affixed to the central axis of the watch [9]. The Rolex design, which evolved from Harwood's original design, remains in use today by many high-end watch manufacturers in the construction of automatic wristwatches. Given the effectiveness of this design in providing mechanical energy to watch designs in the stochastic excitation environment of human wrist movement, it seems appropriate to consider the fitness of such a device as an energy harvester for ocean-based devices.

The objective of this chapter is provide a mathematical model and theoretical basis for the study of a tilt-excited gravitational inertial generator. Such a device may effectively be modeled as a horizontally mounted pendulum with a varying tilt angle. This chapter contributes a fully nonlinear model for a simplified version of the system and conducts a phenomenological and parametric study of its behaviors to provide a

basic understanding of its dynamics before undertaking a design optimization study of the device for energy harvesting.

5.1 Introduction

Mounting a pendulum arm to a bearing with a perfectly vertical axis of rotation creates an interesting astable system that is characterized by the absence of a restoring force. The system is neutrally stable over its entire range of motion, and in the absence of damping, it would continue to rotate at constant velocity indefinitely. The damped system simply spins until its energy is expended through dissipation and the system comes to a halt. Introducing a tilt in the rotational bearing relative to the gravity vector introduces a pair of equilibria: a stable equilibrium at the lowest point in the range of motion, and an unstable equilibrium directly opposite the stable equilibrium. As the tilt angle is changed, the stable equilibrium continuously tracks the lowest point in the range of motion, and the device responds by seeking this equilibrium point.

Because of the tendency of the system to settle into a stable equilibrium arising from any subtle change in its orientation relative to the gravitational field, and due to the exceptionally long natural period of oscillations that occur for very slight tilt angles, the horizontal pendulum has found use in a wide range of applications requiring high-precision measurement. Variations on the prototypical horizontal pendulum have been used to determine the density of the Earth [54], in geological tilt measurement [55], and in seismometry [56]. The seismometry community has made the greatest contributions to the modeling and analysis of the sprung horizontal pendulum, with contributions from Graizer [56, 57] forming the bulk of prior work related to the device. Graizer's work considered devices consisting of one horizontal and two vertical pendula oriented such that the combined measurements of the three could be used to resolve six-degree of freedom (rotational and translational) seismic activity.

In reference [57] he studies the effects of ground tilt on the response and accuracy of the device. Consistent with the application of his research to seismometry, he makes the simplifying assumption of linear pendulum behavior, which is justifiable given the presence of a torsional restoring force in the measurement devices. This chapter expands on the sparse literature base covering this device by providing a model for the fully nonlinear system and studying its behaviors.

The nearest comparable physical system to this device is the planar (vertical) pendulum. Indeed, for a static tilt angle other than horizontal, the device is mathematically equivalent to a vertical pendulum. The driven planar pendulum readily displays a wide range of phenomena, including simple harmonic and subharmonic motion [19] as well as complex behavior such as chaos [58]. The parametrically forced pendulum is of particular interest because of the similarity of the equations of motion to those of a rocking horizontal pendulum. Research by Bartuccelli [59] and Scmitt [60] has shown that high frequency excitation of a parametrically excited pendulum will result in oscillations about a non-zero mean angle. Additionally, parametric excitation can be used for dynamic stabilization of otherwise unstable equilibria [61, 59, 62, 63]. Of particular relevance to this work is reference [64], which considered the influence of tilt on the behavior of a horizontally and vertically excited pendulum, which is similar in mathematical form to the system studied herein. Noting differences in experiment and theoretical results in the references listed above, investigations by Mann [64] identified symmetry-breaking bifurcations resulting from sensitivity of the pendulum's tilt relative to the excitation source.

This chapter is organized as follows: first, a mathematical model for the horizontal pendulum system is derived from first principles. The equations are simulated numerically and example behaviors are shown for the case of harmonic forcing. Three representative behaviors are shown: rotation, oscillation-rotation, and chaos. Of the three, rotating behaviors are the most energetic and thus the most favorable in an

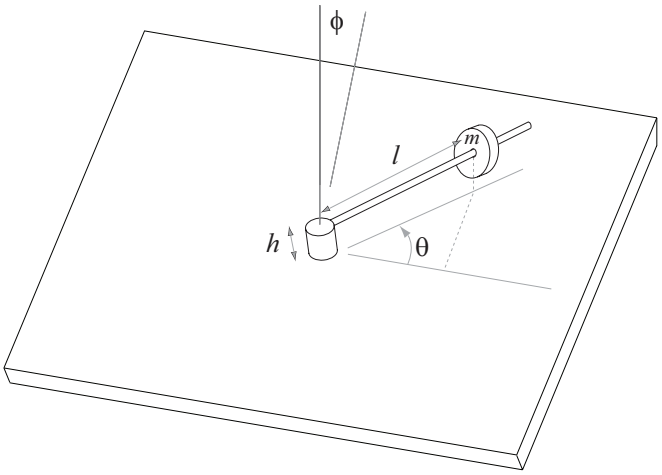

FIGURE 5.1: Schematic of the rocking horizontal pendulum.

energy harvesting context. To further study the rotating solution, a parametric study is conducted using psuedospectral collocation and numerical continuation methods. Conclusions are drawn regarding the basic properties of the device and the information gained is used in the subsequent chapter to guide the design analysis for the device in an ocean environment.

5.2 Physical system and mathematical model

In this section, a mathematical model of the horizontal pendulum system is derived from its physical attributes using Lagrange's equation. In order to maintain a reasonable number of parameters in the physical model, for the purpose of this study the motion of the platform is restricted to tilt about the horizontal axis only.

Figure 5.1 shows a pendulum positioned a vertical distance h above a tilted platform; the angular orientation of the platform is given by the angle x and the angle of the platform is given by ϕ. For a rod of mass m with radius of gyration l measured from the axis of rotation about which x is defined, the kinetic and potential

energies are

$$\mathcal{T} = \frac{1}{2}m\left[\left(l\dot{x}\sin x + h\dot{\phi}\right)^2 + \left(l\dot{\phi}\cos x\right)^2 + \left(l\dot{x}\cos x\right)^2\right] \tag{5.1a}$$

and

$$\mathcal{U} = mg\left[h\cos\phi + l\left(1 - \sin\phi\cos x\right)\right], \tag{5.1b}$$

respectively. It is assumed that $\phi(t)$ is a known, time-varying function and not a generalized coordinate; thus by applying Lagrange's equation for the generalized coordinate x,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial\mathcal{T}}{\partial\dot{x}}\right) - \frac{\partial\mathcal{T}}{\partial x} + \frac{\partial\mathcal{U}}{\partial x} = -b\dot{x}, \tag{5.2}$$

the equation of motion

$$\ddot{x} + \mu\dot{x} + \left(\dot{\phi}^2\cos x + d\ddot{\phi} + \omega^2\sin\phi\right)\sin x = 0 \tag{5.3}$$

is found, where the nonconservative force $Q_x = -b\dot{x}$ corresponds to the viscous damping effects at the bearing, $\mu = \frac{b}{ml^2}$, $d = \frac{h}{l}$ and $\omega^2 = \frac{g}{l}$. Equation (5.3) is similar to the equation for a vertical parametrically excited pendulum, although it lacks a constant restoring force term. The complicated restoring force presents difficulties in studying this model with approximate analytical methods; hence, empirical observations are made using numerical integration to study the behaviors of the system.

5.3 Phenomenology

The previous section provides a model for the horizontal pendulum with an arbitrary, known tilt angle $\phi(t)$. This section conducts a phenomenological investigation into

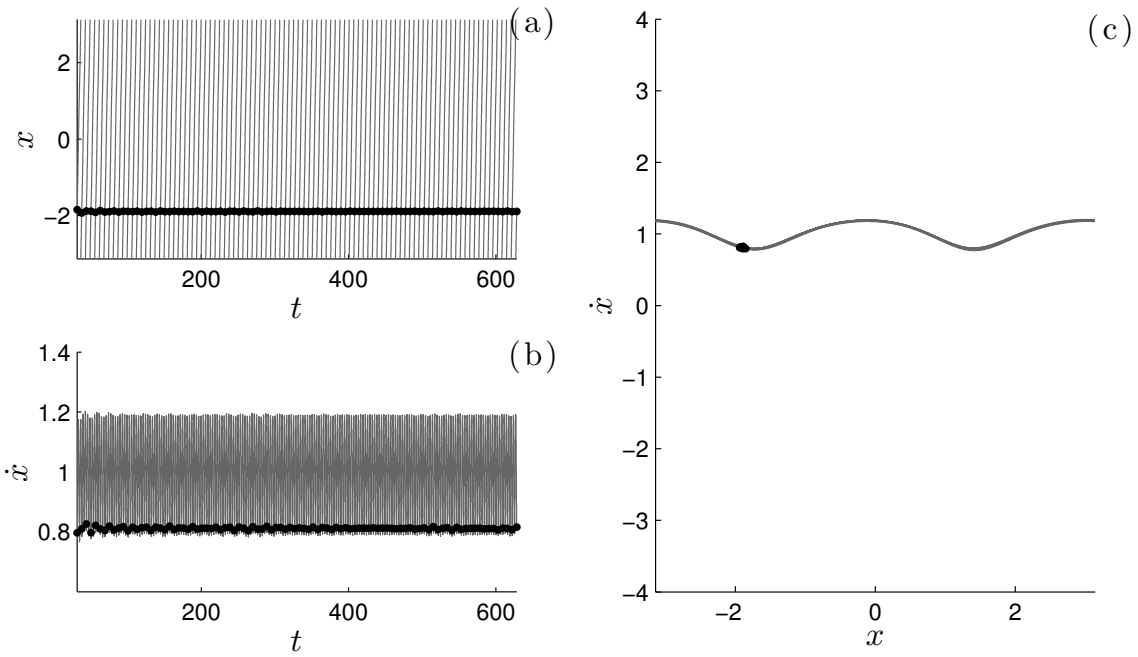

FIGURE 5.2: A stable rotating orbit for parameters $\mu = 0.05$, $d = 0.2$, $\omega = 1$, $\Gamma = 0.7$, and $\Omega = 1$. Shown are the (a) angle, (b) angular velocity and (c) phase portrait of the response for the initial condition $(x_0, \dot{x}_0) = (-1.7903, 0.7592)$. The response is shown as a light gray line and the stroboscopic sections are shown as black dots.

the behaviors of the harmonically forced system with tilt angle $\phi = \Gamma \sin \Omega t$. Given ϕ, the system is now

$$\ddot{x} + \mu x + \left(\Gamma^2 \Omega^2 \cos^2 \Omega t \cos x - d\Gamma\Omega^2 \sin \Omega t + \omega^2 \sin(\Gamma \sin \Omega t)\right) \sin x = 0. \tag{5.4}$$

To explore the behaviors of the system, the state-space equations are numerically integrated using MATLAB's `ode45` integrator [18]. Stroboscopic sections are taken once per forcing period $T_f = \frac{2\pi}{\Omega}$ to determine the periodicity of a solution. These stroboscopic sections form a Poincaré map in the phase space, and by observing the attractor of the Poincaré map, the periodicity of a solution can be classified.

5.3.1 Numerically integrated trajectories

This subsection considers the representative behaviors of equation (5.3) as observed through numerical integration. The harmonically-forced system exhibits three principal behaviors: periodic rotation, periodic oscillation-rotation, and chaotic oscillation-rotation. Figure 5.2 shows an example of a stable rotating orbit simulated for 100 forcing periods. From the figure, it is evident that the orbit is periodic because only one "dot" appears from the stroboscopic section. In this case, the orbit is counterclockwise, as the sign of $\dot{x}$ is positive. There also exists a symmetric counterclockwise orbit due to the symmetry of the system. Stable rotating orbits exist for $\Gamma \approx [0.7, 1.2]$ near $\Omega = 1$ for $\mu = 0.05$, $d = 0.2$ and $\omega = 1$. Investigating the nearby parameter space reveals that a sufficiently large Γ is the largest deciding factor in the presence of stable rotating orbits. In the energy harvesting context, rotating orbits are the most energetically favorable, so it makes sense to design a system that satisfies these conditions.

Figure 5.3 highlights a second type of behavior often encountered in this system. Shown is a periodic oscillating-rotating solution, simulated for 100 forcing periods. Again, the solution is shown to be periodic by the stroboscopic section. Periodic oscillating-rotating solutions often manifest for low forcing frequencies and amplitudes. This behavior contains a superharmonic, as seen by the pinched shape of the phase portrait. Oscillation-rotation occurs when the system swings through one rotation, reverses direction, and swings back in the opposite direction.

Figure 5.4 shows a very common behavior for high-amplitude low-frequency forcing. 100 forcing periods of a chaotic trajectory are shown. The stroboscopic section reveals the chaotic nature of the response, and the trajectory highlights the structure of the chaotic attractor. While the time series for x and $\dot{x}$ appear completely random, the phase portrait $\{x, \dot{x}\}$ is highly structured. Indeed, when the system is simulated

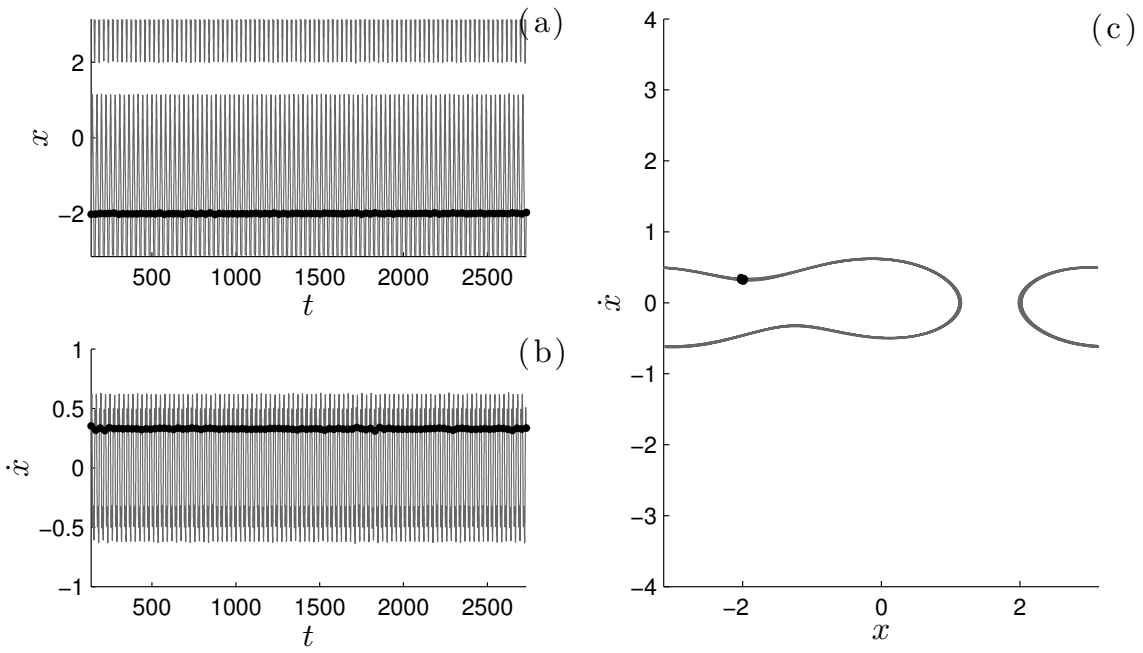

FIGURE 5.3: An oscillating-rotating orbit for parameters $\mu = 0.05$, $d = 0.20$, $\omega = 1.00$, $\Gamma = 0.30$, $\Omega = 0.23$. Shown are the (a) angle, (b) angular velocity and (c) phase portrait of the response for the initial condition $(x_0, \dot{x}_0) = (-3.14, 3.00)$. The response is shown as a light gray line and the stroboscopic sections are shown as black dots.

for 10,000 periods, the strange attractor appears quite structured, as seen in figure 5.5. The system tends to oscillate chaotically around the alternating stable equilibria at $x = \{0, \pm n\pi\}$ and occasionally rotate through to the opposite equilibrium point.

5.3.2 Basins of attraction of rotating solutions

The rotating solution shown in figure 5.2 is the primary behavior of interest in this system. In designing an energy harvester with this device, it may prove valuable to exploit clockwise or counterclockwise rotations because of the design of the energy harvester. As the device is symmetric, for every clockwise stable rotating orbit, there also exists a counterclockwise stable rotating orbit. The device will settle on one or the other depending on the choice of initial conditions. This can be demonstrated with a basins of attraction diagram, which scatters initial conditions $\{x_0, \dot{x}_0\}$ over

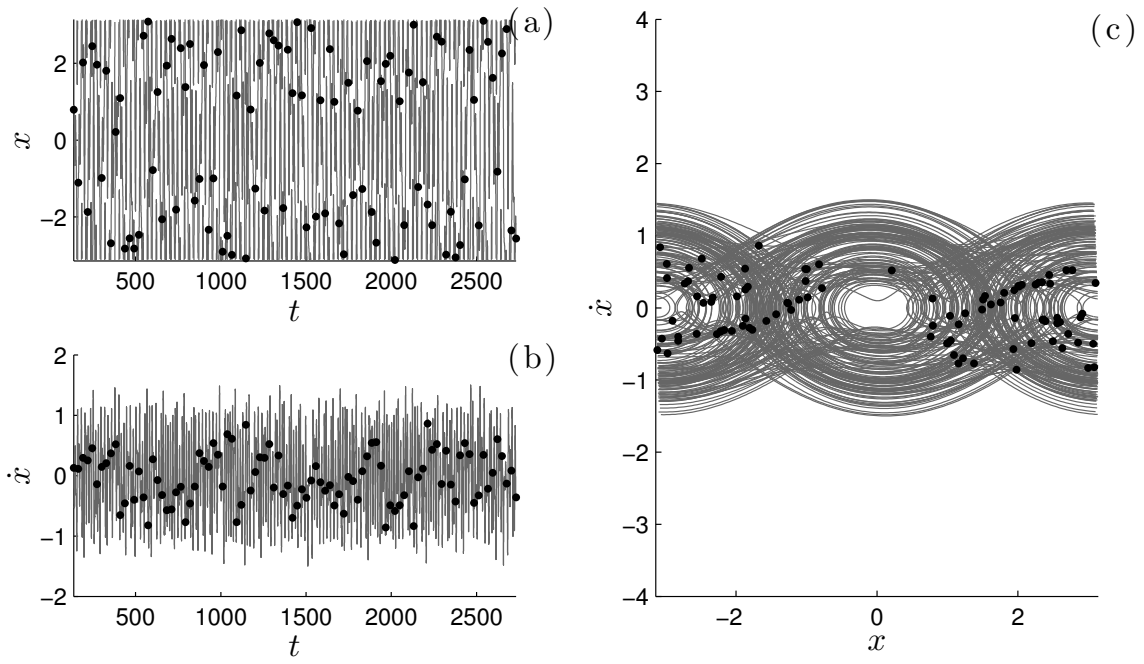

FIGURE 5.4: A chaotic trajectory for parameters $\mu = 0.05$, $d = 0.20$, $\omega = 1.00$, $\Gamma = 0.70$, $\Omega = 0.23$. Shown are the (a) angle, (b) angular velocity and (c) phase portrait of the response for the initial condition $(x_0, \dot{x}_0) = (-3.14, 3.00)$. The response is shown as a light gray line and the stroboscopic sections are shown as black dots.

an interval of interest and classifies each initial condition as precursor to either a clockwise or counterclockwise orbit.

Figure 5.6 shows the basins of attraction for two symmetric rotating solutions for the range of initial conditions $x \in [-\pi, \pi]$ and $\dot{x} \in [-2, 2]$. A 500 by 500 grid of evenly spaced initial conditions was chosen and categorized based on the direction of spin observed (i.e. the sign of $\dot{x}(t_f)$). Clockwise rotations are originate from the initial conditions in the gray regions, while counterclockwise rotations originate from the initial conditions in the white regions. The basin boundaries are fractal, as is typical in a nonlinear system with coexisting solutions, and they illustrate the high sensitivity of the system to initial conditions.

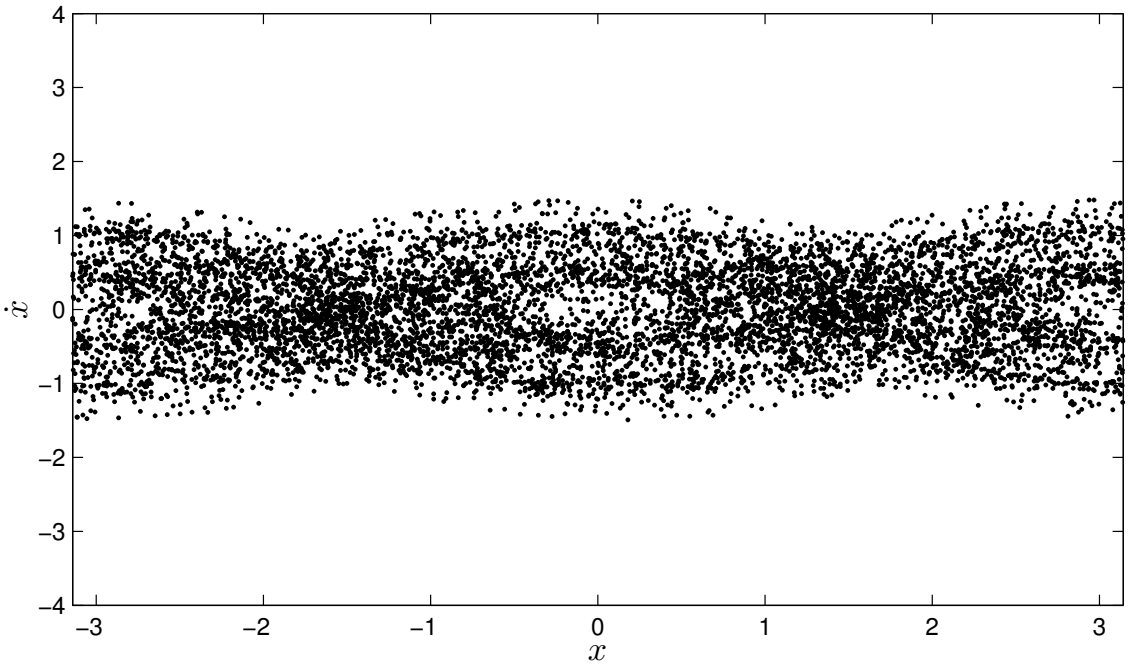

FIGURE 5.5: Strange attractor for parameters $\mu = 0.05$, $d = 0.20$, $\omega = 1.00$, $\Gamma = 0.70$, $\Omega = 0.23$, simulated for 10,000 periods.

5.4 Numerical continuation of periodic rotating orbits

In the previous section, typical behaviors were discussed and representative examples were shown. Of the behaviors shown, the periodic rotating orbits are of the most interest in an energy harvesting context. In order to study this behavior further, psuedospectral collocation is used in conjunction with numerical continuation to follow the orbits as parameters are varied.

The horizontal pendulum system presents a unique challenge for the continuation of rotating orbits with psuedospectral collocation. Because of the periodicity of the angular variable x in addition to the time-periodicity of the system, a Fourier basis, the typical basis of choice when solving a periodic BVP, is no longer useable. Instead, a Chebyshev basis is used and boundary conditions are incorporated that specify the type of orbit sought. This section describes the procedure used to obtain these orbits from the construction of the time-periodic boundary value problem (TP-BVP) to the

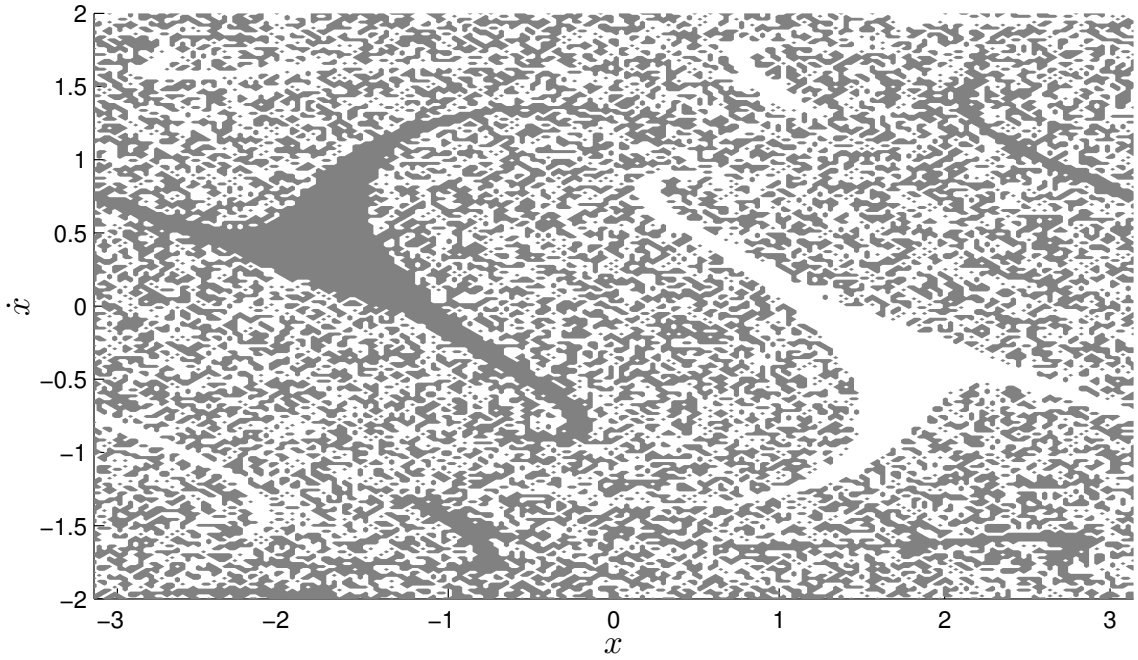

FIGURE 5.6: Basins of attraction in the phase space for stable rotating solutions for parameters $\mu = 0.05$, $\omega = 1$, $\Gamma = 0.7$, $\Omega = 1$. The shaded area corresponds to a clockwise rotating solution, while the unshaded area corresponds to a counter-clockwise rotating solution.

use of numerical continuation to follow the orbits.

5.4.1 Formulation of the TP-BVP

This section formulates the TP-BVP used to compute the periodic orbits of equation (5.3). In order to compute the periodic orbits of equation (5.3), it is necessary to rewrite the equation in state-space form. Equation (5.3) can we rewritten in state-space form as

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \frac{1}{\Omega} \begin{pmatrix} x_2 \\ -\mu x_2 + \left(\dot{\phi}^2 \cos x_1 + d\ddot{\phi} + \omega^2 \sin \phi \right) \sin x_1 \end{pmatrix} = f(\boldsymbol{x}, t; \boldsymbol{p}) \tag{5.5}$$

where $\phi = \Gamma \sin(2\pi t)$ and the factor $\frac{1}{\Omega}$ rescales the equation to the domain $t \in [0, 1]$. The parameter set $\boldsymbol{p}$ includes μ, the viscous damping coefficient, $d = \frac{h}{l}$, the relative

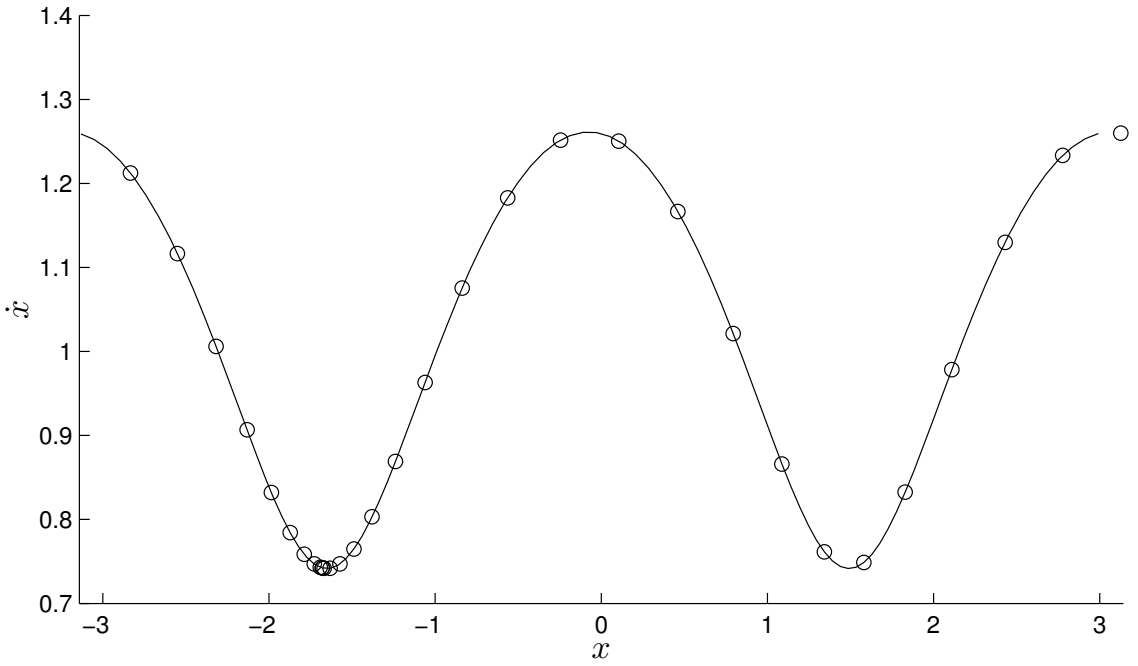

FIGURE 5.7: Example orbit found by collocation (o) overlaid on numerical simulation (-) for $\mu = 0.05$, $d = 0.2$, $\omega = 1$, $\Gamma = 1$ and $\Omega = 1$. The collocation points are clustered near points corresponding to $t = 0$ and $t = 1$.

height above the platform, ω, the linear natural frequency, Γ, the forcing amplitude and Ω the forcing frequency.

To discretize the problem, a Chebyshev differentiation matrix $\boldsymbol{D}$ is generated along with a set of collocation points t_c as described in [49]. Unlike the Fourier differentiation matrix discussed in chapter 4, periodic boundary conditions are not explicitly included in the Chebyshev differentiation matrix. Thus $\boldsymbol{D}$ must be altered as follows to incorporate the boundary conditions:

$$\boldsymbol{D} = \begin{pmatrix} D_{1,1} & D_{1,2} & \cdots & D_{1,n-1} & D_{1,n} \\ D_{2,1} & \ddots & & & \vdots \\ \vdots & & \ddots & & \vdots \\ D_{n-1,1} & D_{n-1,2} & \cdots & D_{n-1,n-1} & D_{n-1,n} \\ 1 & 0 & \cdots & 0 & -1 \end{pmatrix} \tag{5.6}$$

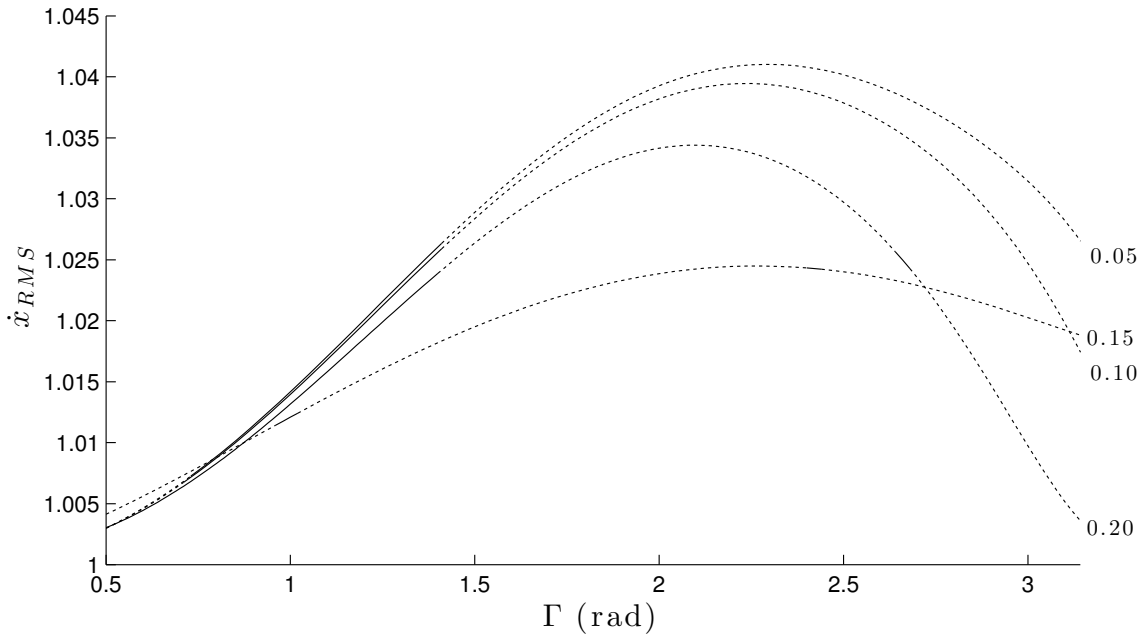

FIGURE 5.8: RMS velocity of rotating solution as the rocking amplitude Γ is varied for parameters $\mu = \{0.05, 0.10, 0.15, 0.20\}$, $d = 0.2$, $\omega = 1$ and $\Omega = 1$. Stable solutions exist on the solid lines, while unstable solutions are found on the dotted lines.

Then, by zeroing the residual

$$R(\boldsymbol{x}(\boldsymbol{t_c}); \boldsymbol{p}) = \boldsymbol{I}_2 \otimes \boldsymbol{D}\boldsymbol{x} - f(\boldsymbol{x}, t; \boldsymbol{p}) \tag{5.7}$$

periodic rotating orbits of the system can be found. Note that for the rows of $\boldsymbol{D}$ corresponding to boundary conditions, $f(\boldsymbol{x}, t; \boldsymbol{p})$ is replaced with the appropriate boundary condition: $x(0) - x(1) = \pm 2\pi$ (depending on the direction of rotation desired) and $\dot{x}(0) - \dot{x}(1) = 0$. This residual is then continued using the methods previously described in this dissertation.

5.4.2 Response trends of the rotating orbits

By solving the TP-BVP presented in the previous subsection and continuing the solution as parameters are varied, the effects of changing model parameters on the existence and stability of rotating orbits can be observed.

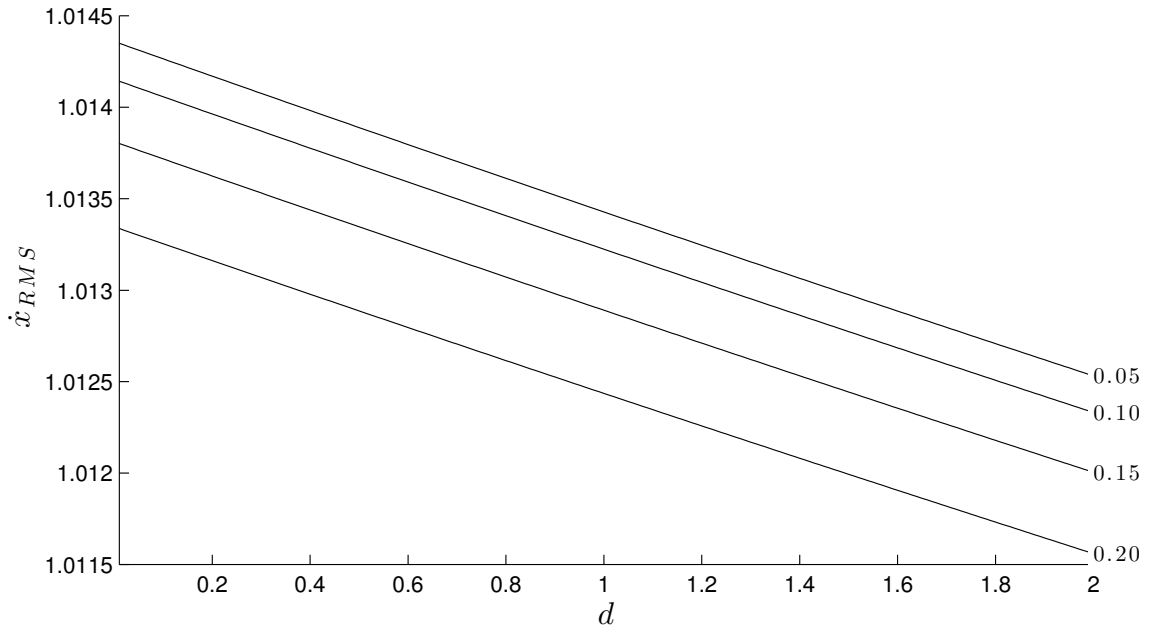

FIGURE 5.9: RMS velocity of rotating solution as the arm height ratio d is varied for parameters $\mu = \{0.05, 0.10, 0.15, 0.20\}$, $\omega = 1$, $\Gamma = 1$ and $\Omega = 1$. Stable solutions exist on the solid lines, while unstable solutions are found on the dotted lines.

From this investigation, it is observed that the parameters with the most notable effect on the response are μ and Γ. The rocking amplitude Γ seems to have the strongest effect on the response. It is observed that stable rotating orbits only exist for moderate ranges of Γ. Figure 5.8 shows the influence of Γ on the root-mean-square velocity of periodic rotating orbits. It can be seen that the peak predicted orbit is unstable, and is also affected by the damping parameter μ. As μ increases, the orbits generally trend downwards, but the trend is not linear (c.f. $\mu = \{0.15, 0.20\}$).

Figure 5.9 shows the influence of arm height above the platform. As the arm is raised above the platform (and thus $d = h/l$ increases), the response trends downward linearly. In this case, the influence of the damping parameter μ is again nonlinear, as can be seen by the spacing of the lines in figure 5.9.

5.5 Conclusions

From the investigations conducted in this chapter, it is apparent that the most desirable behavior for energy harvesting exhibited by this system is rotation. The rotating solutions of the system exist only in a limited region of the parameter space, as can be seen by the numerical continuation results. Of the parameters studied, the tilt angle Γ appears to have the greatest effect on the existence of rotating solutions, provided that the linear natural frequency ω is somewhat commensurate with the forcing frequency Ω. In tuning ω, it makes the most sense to ensure that $\omega \approx \Omega$, but to err on the side of $\omega < \Omega$. Given the sensitivity of the rotating orbits to changes in the parameters, it is unlikely that a purely rotating orbit will manifest in the stochastically forced system to be considered in the following chapter. However, solutions with a predominately rotating component may appear, and such solutions would be the most favorable from an energy harvesting perspective.

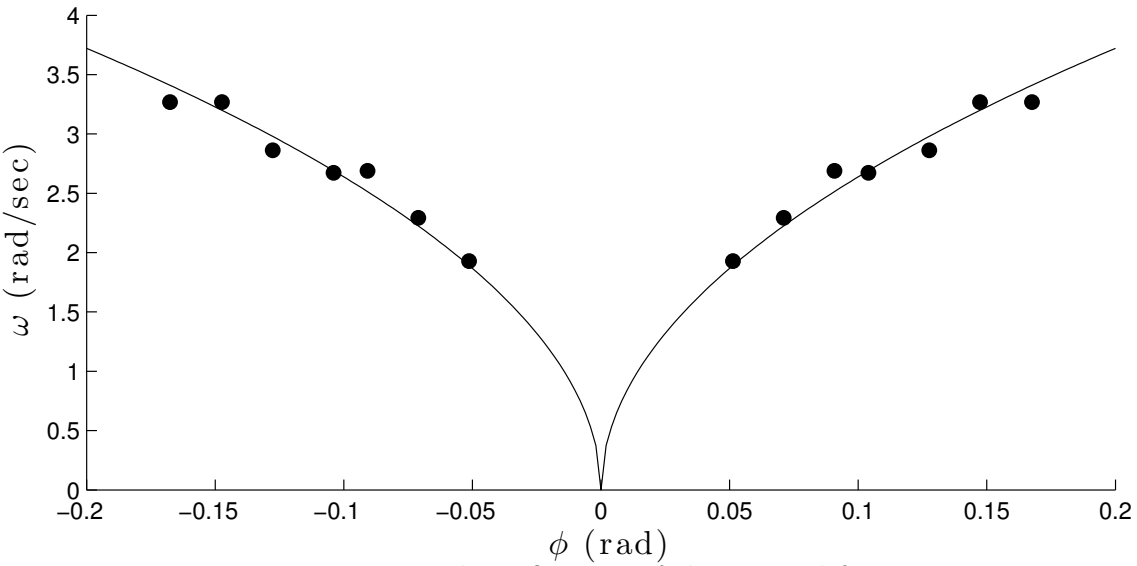

FIGURE 5.10: Experimental verification of the natural frequency curve.

Supplement: Experimental parameter measurement and identification

Early in the process of investigating this device, an experiment was constructed with the intent of conducting a range of experimental tests to confirm the model shown in equation (5.3). While the device proved difficult to control and was unable to conduct extensive verification of the model presented in this chapter, we were able to verify the natural frequency curve for various tilt angles. The linear natural frequency for a tilted pendulum is given by

$$\omega_n^2 = \frac{g}{l} \sin \phi. \tag{5.8}$$

In figure 5.10, experimentally measured tilt angles and natural frequencies are compared with the predictions of equation (5.8). By using MATLAB's `fminunc`, a nonlinear optimization routine for unconstrained objective functions, the radius of gyration l for the experimental system was determined to be 0.1407 m, providing the fit seen in figure 5.10.

6

Design of an Ocean Energy Harvester

Designing an energy harvester that functions effectively in an ocean environment requires consideration of the stochastic effects introduced by ocean wave forcing. The tilt-excited gravitational inertial generator device proposed in the previous chapter exhibits high-energy rotating behavior that is desirable for energy harvesting, but in the absence of an active or passive control system, it is unlikely that a stable rotating orbit will manifest in the random ocean excitation environment. Rather, the device will need to function in much the same way as the self-winding wristwatch mechanisms, by providing spurts of energy to an energy storage medium at random intervals when a sufficiently energetic behavior manifests itself. Thus, optimizing the design of such a device should focus more on the long-term availability of kinetic energy that could be converted to and stored as electrical energy, rather than on maintaining any one type of behavior in particular.

The objective of this chapter is to propose a set of design recommendations for the construction of a prototype device to be used in future in-situ experiments. In order to determine an effective design strategy for the ocean energy harvesting system, it is necessary to collect a large amount of empirical data and determine a design strategy

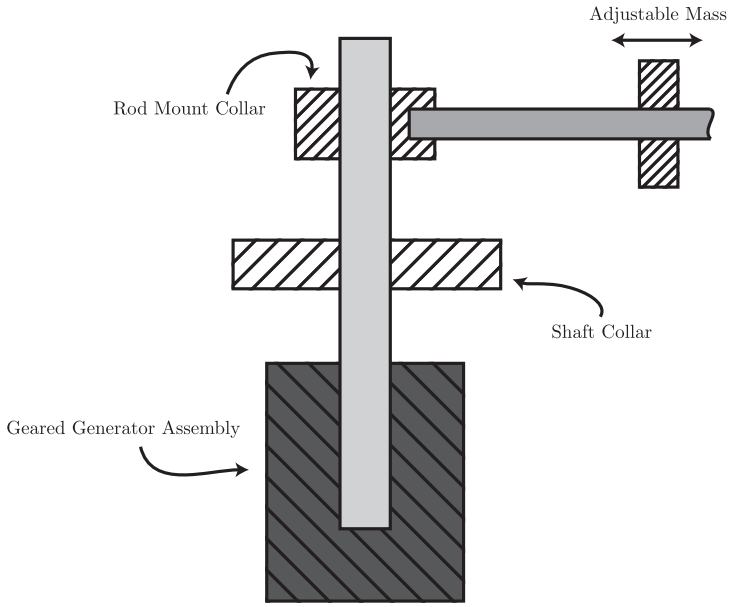

FIGURE 6.1: Conceptual drawing of the gravitational inertial generator with key components labeled. The device includes an adjustable mass for tuning.

from empirical observations. This chapter employs a computational model based on the physical model derived in the previous chapter to conduct numerical experiments on the device over a reasonable range of parameters and sea states. From the data collected, a response surface is constructed that can be used to optimize the system and guide future design decisions when constructing a physical device.

6.1 Introduction

In creating design recommendations for a prototype of the energy harvesting system shown in figure 6.1, it is necessary to step beyond deterministic analyses of the device as a benchtop experiment and consider how it will behave in its target environment. By employing the lessons learned from the deterministic investigations discussed in the previous chapter, this chapter conducts a design study for a range of sea states and design parameters.

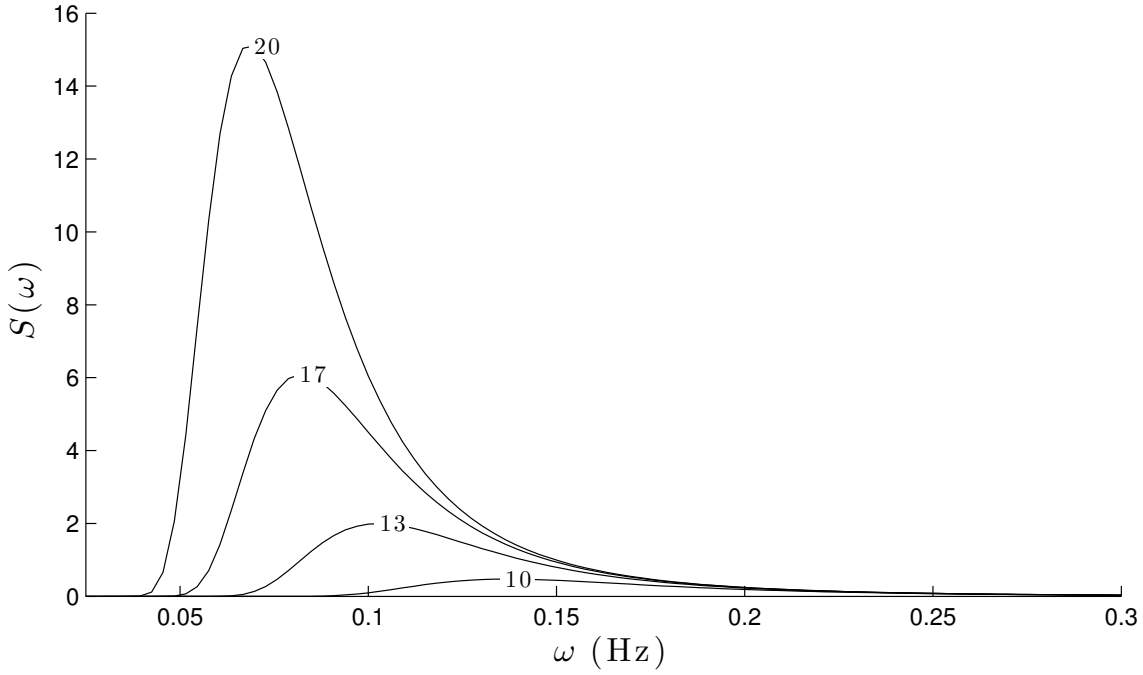

FIGURE 6.2: Pierson-Moskowitz wave spectra for wind speeds $U_{19.5} = \{10, 13, 17, 20\}$ m/s.

Ocean waves are generated by many physical processes, including convection from thermal energy absorbed from solar radiation, gravitational effects from the moon, and the interaction of the atmosphere with the water surface in wind currents. When the sea reaches a state of equilibrium where the energy added by wind currents is matched by the energy transported in wave motion and dissipated in wave breaking, the sea is said to be fully developed. In a fully developed sea, the power spectral density (PSD) of the waves is stationary, and representations of PSDs for fully developed seas are frequently used in marine engineering and oceanography. A frequently used model of the PSD of fully developed seas arose from the work of Willard Pierson and Lionel Moskowitz [1]. By observing sea states and wind speeds using weather ships, they were able to construct an empirical PSD of ocean waves for various wind speeds and fit an analytical model to their observations. Their model, now known as the Pierson-Moskowitz spectrum, is used in this investigation to simulate random

ocean waves.

This chapter is organized as follows: first, a discussion of ocean excitation environments is presented to familiarize the reader with the process of modeling randomly generated ocean waves. The Pierson-Moskowitz spectrum is used in conjunction with the Bennett-Rice approximate formulation of a stochastic process to generate realizations of the surface elevation and slope of ocean waves. The model presented in the previous chapter is then reviewed in the context of random forcing and examples of randomly forced behavior are shown. A series of experiments are designed using the Sobol sequence [65], and numerical experiments are conducted from this design. Finally, the results of the numerical experiments are fit by a response surface reconstruction and conclusions are drawn regarding the effective design of an energy harvesting device with the horizontal pendulum system as the method of kinetic energy capture.

6.2 Power spectra of ocean waves

In this section, the Pierson-Moskowitz spectrum is presented along with the details of simulating random waves.

6.2.1 The Pierson-Moskowitz spectrum

The Pierson-Moskowitz spectrum is given by

$$S(\omega) = \frac{\alpha g^2}{\omega^5} \exp\left[-\beta \left(\frac{\omega_0}{\omega}\right)^4\right], \tag{6.1}$$

where $\omega = 2\pi f$ is the wave frequency in radians per second (f is the wave frequency in Hertz), $\alpha = 8.1 \times 10^{-3}$, $\beta = 0.74$, $\omega_0 = g/U_{19.5}$ and $U_{19.5}$ is the windspeed at 19.5 meters[1] above the surface [1]. A plot of the Pierson-Moskowitz spectrum for various

[1] This was the height of the anemometers on the weather ships used in their investigation [1].

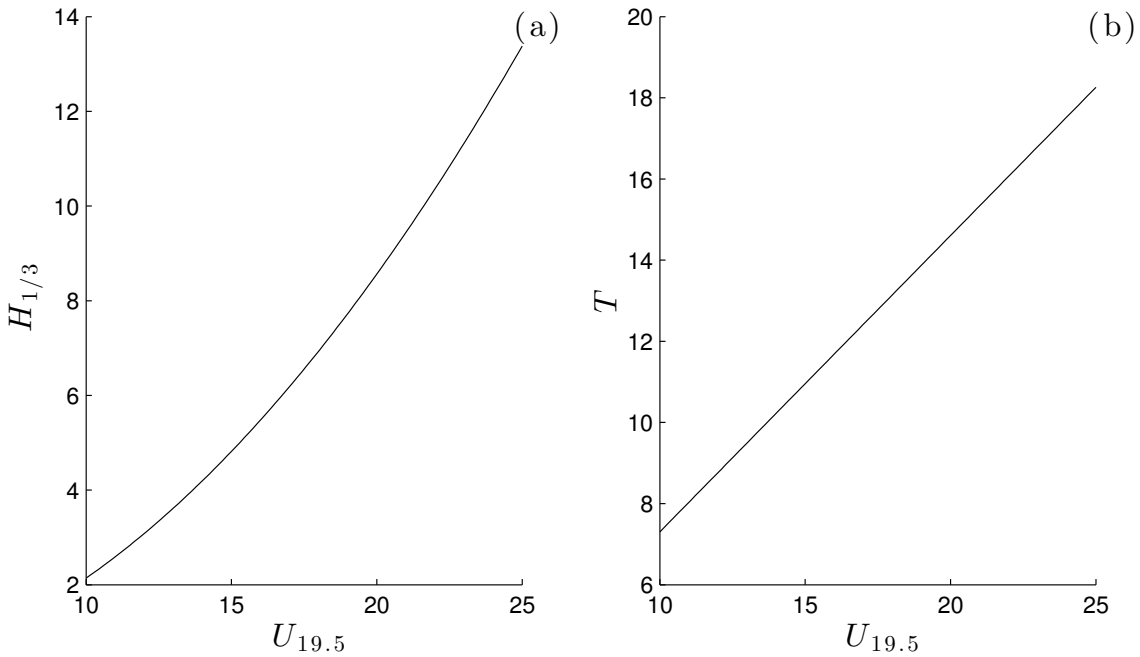

FIGURE 6.3: Significant wave height (a) and period (b) of deep-water waves in fully developed seas as a function of wind velocity 19.5 meters above the surface $U_{19.5}$ as determined by the Pierson-Moskowitz spectrum.

windspeeds $U_{19.5}$ is shown in figure 6.2.

Deep-water ocean waves in fully developed seas are characterized by long periods and significant wave heights ranging from a few meters to the height of a four-story building [1, 66], as can be seen in figure 6.3. Other interesting properties of the spectrum can be observed by manipulating equation (6.1). By solving $\partial S_{PM}(\omega)/\partial\omega = 0$, the peak frequency is found to be

$$\omega_p = \frac{0.877g}{U_{19.5}}. \tag{6.2}$$

From this information, the wave speed c_p of the waves at the peak frequency can be determined as

$$c_p = \frac{g}{\omega_p} = 1.14U_{19.5}. \tag{6.3}$$

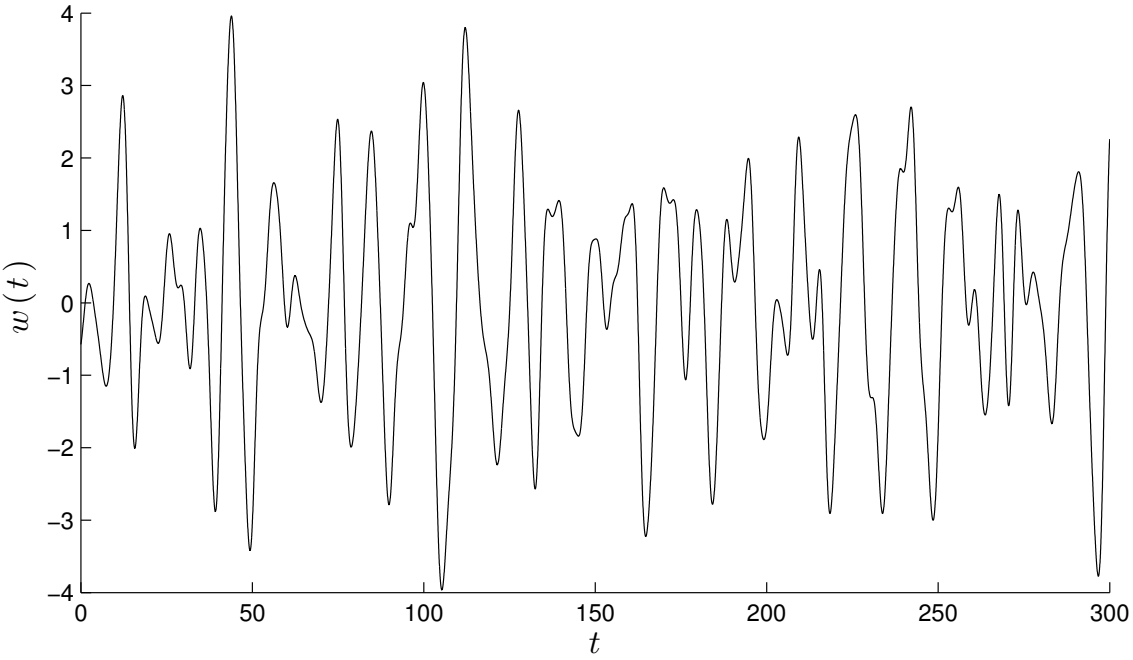

FIGURE 6.4: One realization of a wave-height $w(t)$ time history generated from the Pierson-Moskowitz spectrum for $U_{19.5} = 20$ m/s.

In addition to describing the frequency content of ocean waves, the Pierson-Moskowitz spectrum implicitly describes the spatial attributes of the waves through the dispersion relation [66]. The dispersion relation for deep water waves relates the frequency of the wave to its wavelength as

$$\lambda = \frac{2\pi g}{\omega^2}. \tag{6.4}$$

This information can now be used to generate random waves from the Pierson-Moskowitz spectrum.

6.2.2 Generating random waves

To generate a random wave, it is assumed that the power spectrum can be divided into N frequency components ω_i and reconstructed as a series of sinusoids with random phase angles. Under these assumptions, the surface elevation $\eta(x,t)$ can be

written

$$\eta(x,t) = \sum_{i=1}^{N} a_i \sin(k_i x - \omega_i t + \psi_i) \tag{6.5}$$

where $k_i = \frac{2\pi}{\lambda_i}$, $a_i = [2S(\omega_k)\Delta\omega]^{1/2}$, and the random phase angle ψ_i is drawn from a uniform distribution over $[0, 2\pi]$. Figure 6.2 shows one realization of a Pierson-Moskowitz process as realized through this approximation, measured at a single location.

For the energy harvester design considered in this chapter, the wave height does not appear in the equation. Rather, the wave slope angle $\phi(x,t)$ and its temporal derivatives appear in the equation as forcing terms. To find the wave slope, $\eta(x,t)$ is differentiated with respect to x, yielding

$$\frac{\partial \eta(x,t)}{\partial x} = \tan[\phi(x,t)] = \sum_{i=1}^{N} a_i k_i \cos(k_i x - \omega_i t + \psi_i). \tag{6.6}$$

To simplify the model, it is assumed that the buoy does not drift from a single location x. Given this assumption, substituting

$$k_i = \frac{2\pi}{\lambda_i} = \frac{\omega_i^2}{g} \tag{6.7}$$

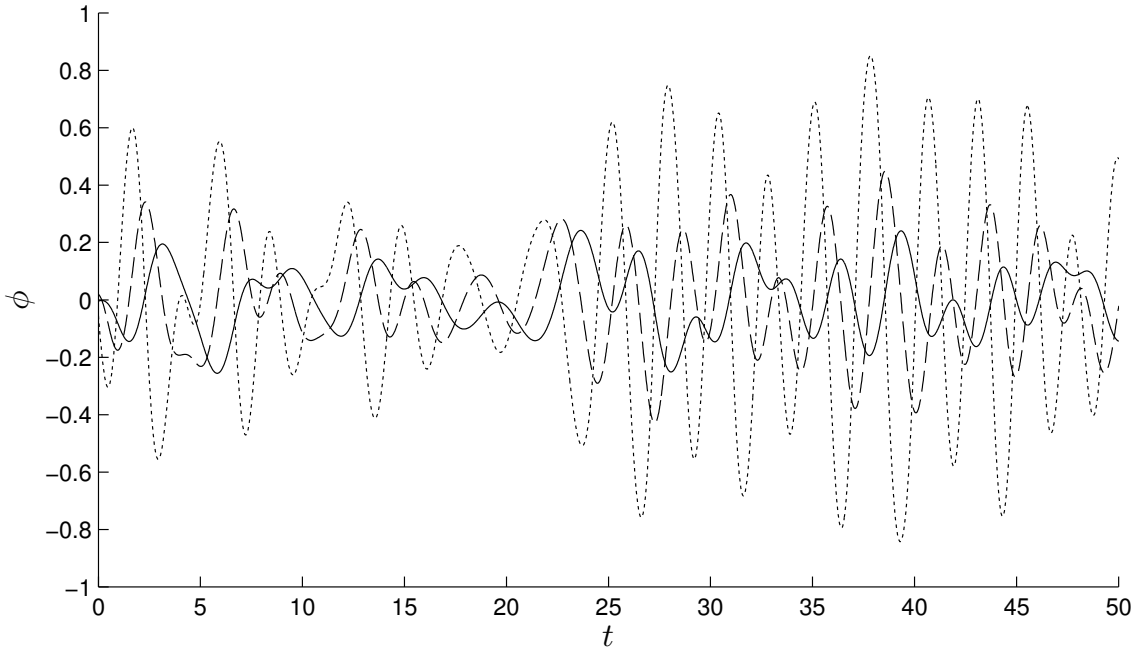

FIGURE 6.5: One realization of a time history for ϕ (solid line) and its derivatives $\dot{\phi}$ (dashed line) and $\ddot{\phi}$ (dotted line) generated from the Pierson-Moskowitz spectrum for $U_{19.5} = 20$ m/s using the formulae given in equations (6.8).

yields

$$\phi(t) = \arctan\left[\sum_{i=1}^{N} a_i \left(\frac{\omega_i^2}{g}\right) \sin(\omega_i t + \psi_i)\right] \tag{6.8a}$$

$$\dot{\phi}(t) = \frac{\sum_{i=1}^{N} a_i \left(\frac{\omega_i^3}{g}\right) \cos(\omega_i t + \psi_i)}{1 + \left(\sum_{i=1}^{N} a_i \left(\frac{\omega_i^2}{g}\right) \sin(\omega_i t + \psi_i)\right)^2} \tag{6.8b}$$

$$\begin{aligned}\ddot{\phi}(t) = &-\frac{\sum_{i=1}^{N} a_i \left(\frac{\omega_i^4}{g}\right) \sin(\omega_i t + \psi_i)}{1 + \left(\sum_{i=1}^{N} a_i \left(\frac{\omega_i^2}{g}\right) \sin(\omega_i t + \psi_i)\right)^2} \\ &- \frac{2\left(\sum_{i=1}^{N} a_i \left(\frac{\omega_i^3}{g}\right) \cos(\omega_i t + \psi_i)\right)^2 \left(\sum_{i=1}^{N} a_i \left(\frac{\omega_i^2}{g}\right) \sin(\omega_i t + \psi_i)\right)}{\left(1 + \left(\sum_{i=1}^{N} a_i \left(\frac{\omega_i^2}{g}\right) \sin(\omega_i t + \psi_i)\right)^2\right)^2}.\end{aligned} \tag{6.8c}$$

Note that due to assumption of a fixed spatial location and the arbitrary phase angles ψ_i, cosine may be replaced with sine and x may be eliminated without loss of generality. These expressions may now be substituted into the model for the system and used in numerical simulations.

6.3 Mathematical model and numerical considerations

The previous section described a method for generating random ocean waves. The purpose of this section is to discuss how to apply those results to the model derived in the previous chapter. From chapter 5, it is know that the equation of motion for a horizontal pendulum with tilt angle ϕ is given by

$$\ddot{x} + \frac{b}{ml^2} + \left[\frac{h}{l}\ddot{\phi} + \dot{\phi}^2 \cos x + \frac{g}{l}\sin\phi\right]\sin x = 0. \tag{6.9}$$

Chapter 5 considered the case $\phi = \Gamma \cos \Omega t$, whereas the present chapter considers ϕ to be defined as in equation (6.8). To study the effects of random forcing on equation (6.9), the equation may be numerically integrated.

Figure 6.6 shows an example of a response for randomly generated waves from a Pierson-Moskowitz spectrum. The device exhibits chaotic behavior and does not lock into a steady-state solution because of the random excitation. This is uniformly the case for random excitation of this device. Thus it no longer makes sense to consider the Poincaré section or similar metrics in classifying the device's behavior. Instead, root mean square velocity

$$\dot{x}_{RMS} = \sqrt{\frac{1}{T_f}\int_0^{T_f} \dot{x}^2(t)\,\mathrm{d}t} \tag{6.10}$$

is used to delineate between the performances of the different simulations.

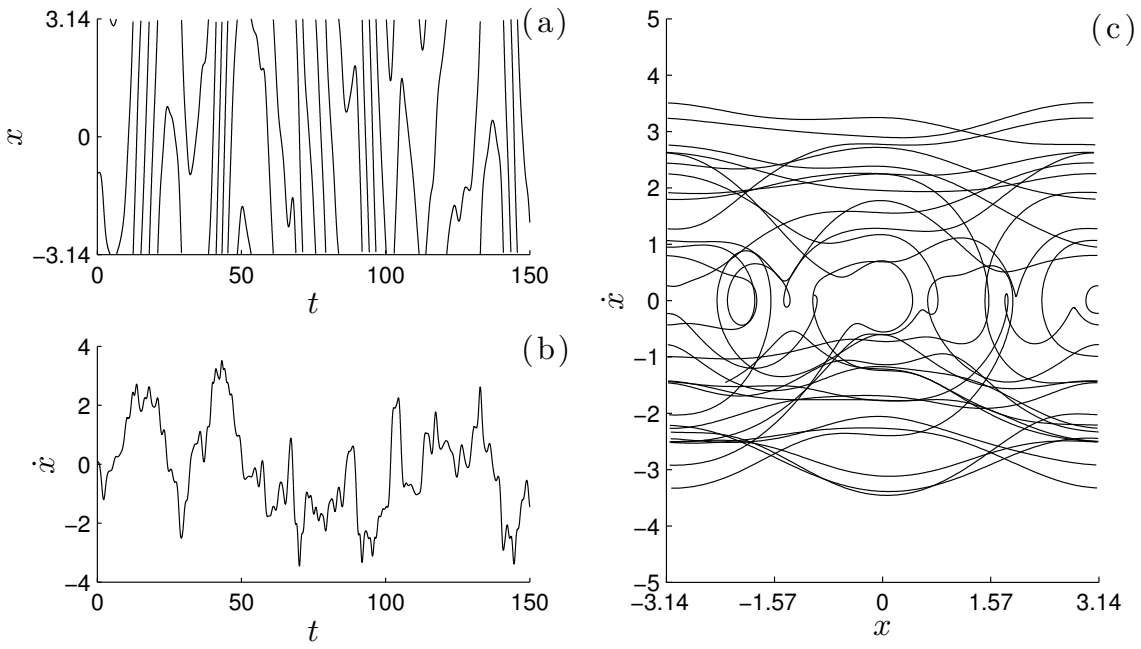

FIGURE 6.6: Response of the device with Pierson-Moskowitz wave forcing for $U_{19.5} = 20$ m/s. Other parameters are $\frac{b}{m} = 0.05$, $h = 0.2$ and $l = 1$.

6.3.1 Numerical considerations

Accomplishing the primary objective of this investigation, reconstructing the response surface of equation (6.9), requires many thousands of simulations scattered across the parameter space to determine a reliable estimate for the surface. Although equation (6.9) is relatively straightforward to numerically integrate, there are a few items to consider with regards to code optimization given the fact that many thousands of simulations will be performed. Three major concerns are presented in this subsection: the length of simulation needed to accurately estimate $\dot{x}_{RMS}$; efficiently generating ϕ and its derivatives; and sampling $\phi(t)$ at arbitrary time intervals.

In order to construct an accurate representation of the response surface, it is necessary to ensure that the RMS velocity of equation (6.9) is accurately observed at each iteration. This requires the equation to be simulated for a long period of time to ensure transient behavior does not influence the results. Figure 6.7 shows the

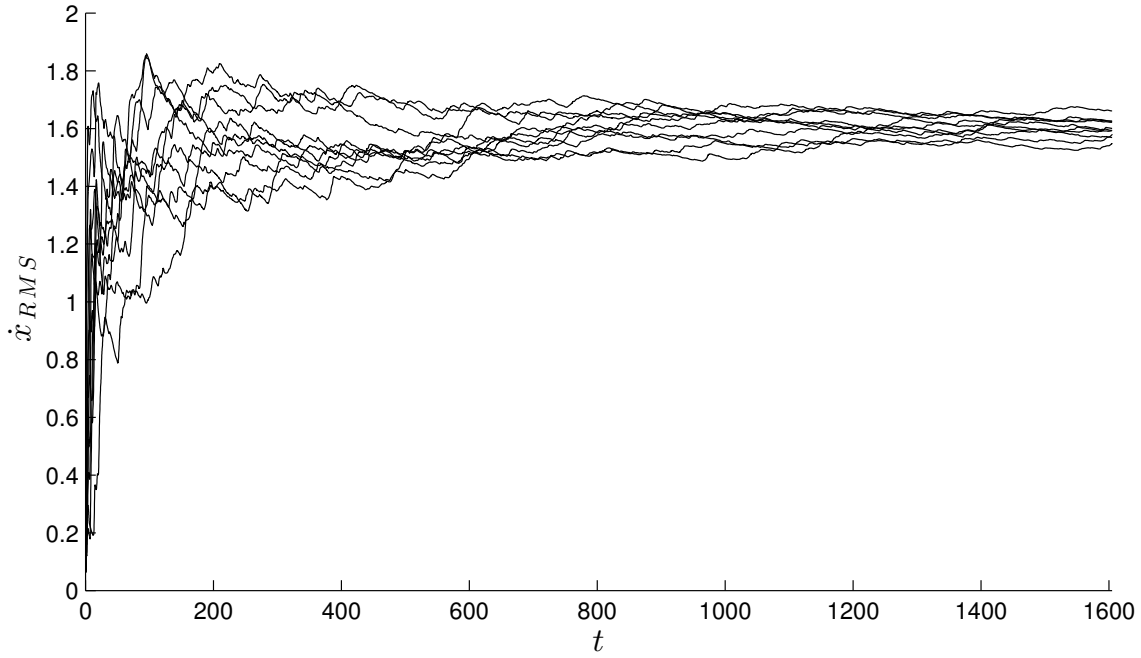

FIGURE 6.7: The RMS velocity slowly converges to a stable value as the length of the simulation is increased. Shown here are ten simulations for $l = 1$, $\xi = 0.05$ and $h = 0.2$.

trend in measured RMS velocity of equation (6.9) simulated for several realizations of a Pierson-Moskowitz wave history. As the length of the simulation increases, all realizations converge to the same RMS velocity. Experimentation yielded an effective simulation length of 800 periods, i.e.

$$T_s = 800\frac{2\pi}{\omega} \tag{6.11}$$

to ensure convergence of the estimate for $\dot{x}_{RMS}$.

For each iteration of the numerical experiment presented in the following section, a new realization of ϕ and its derivatives must be computed from equations (6.8). In order to complete this task efficiently, it is advantageous to eliminate the summation and utilize MATLAB's strong linear algebra capabilities. MATLAB is built on the Intel Math Kernel Library (IMKL), and as such, replacing sum-

mations with dot products allows MATLAB to access the highly optimized IMKL Basic Linear Algebra Subprograms (BLAS) library. For a frequency range of N sample frequencies $\boldsymbol{\omega} = \{\omega_1, \omega_2, \ldots, \omega_N\}^\top \in \mathbb{R}^{N\times 1}$, corresponding amplitudes $\boldsymbol{a} = \{a_1(\omega_1), a_2(\omega_2), \ldots, a_N(\omega_N)\}^\top \in \mathbb{R}^{N\times 1}$, and a time span of T times $\boldsymbol{t} = \{t_1, t_2, \ldots, t_T\} \in \mathbb{R}^{1\times T}$, equations (6.8) can be rewritten as

```
phi = atan((1/g)*(w.*w.*a)'*sin(w*t + psi));
```

and so forth (rather than using the `sum` command). Precomputing ϕ and its derivatives rather than using an explicit formula during numerical integration of equation (6.9) often reduces computational time by several seconds, an important savings considering the number of runs required to reconstruct the response surface.

To efficiently sample ϕ and its derivatives, simple linear interpolation over the precomputed wave realization is simple and effective. MATLAB's `interp1q` utility performs the interpolation quickly (without error checking) and is useful given the large number of runs to be performed in the following investigation. By providing a way to sample ϕ and its derivatives at arbitrary values of t, the integrator is able to take adaptive step sizes, greatly reducing computational time.

6.4 Design of a numerical experiment for response surface reconstruction

The previous section reviewed the model derived in chapter 5 and provided an example of the response for one realization of a Pierson-Moskowitz wave history. The focus of this section is to construct the response surface over a range of wind speeds and design parameters to provide a basis for design decisions.

6.4.1 Quasi-Monte Carlo sampling with the Sobol sequence

Given the random nature of the ocean wave forcing and the nonlinearities in the response of the device, parametric studies such as those in the previous few chapters

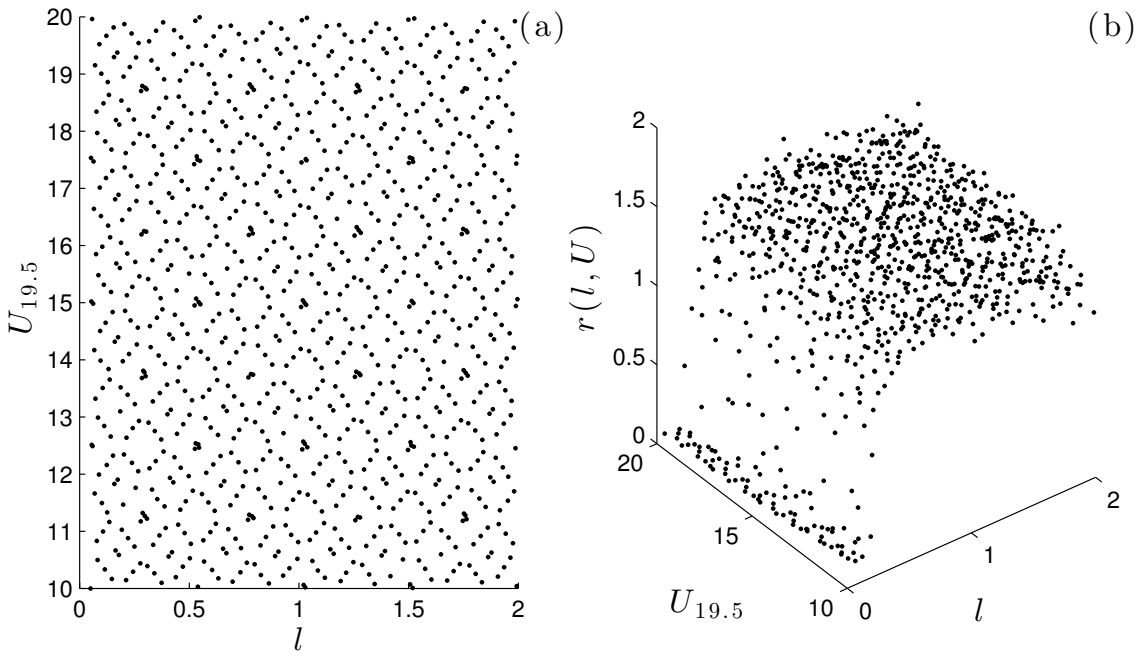

FIGURE 6.8: Results of a qMC experiment $\mathcal{E}(1000, 2)$ for the horizontal pendulum system. Figure (a) shows the sample coverage of the $\{l, U\}$ space, while figure (b) shows the response of samples projected into the $\{l, U\}$ space.

are of little use in predicting the response of the device. In this section, a quasi-Monte Carlo (qMC) investigation of the parameter space is conducted using the Sobol sequence to ensure an accurate representation of the response surface [65]. Utilizing qMC sampling as opposed to purely random MC sampling or latin hypercube sampling ensures better coverage of the sample space in fewer iterations than would be required with a true MC method.

Denote by $\mathcal{E}(N, P)$ the qMC experiment with N samples in a space of P parameters. Each parameter P_i is formally assumed to be uniformly distributed, but in practice each realization of P_i is drawn from the Sobol sequence scaled to its range, i.e. $P_{n,i} \sim \text{Sobol}_{(n)}([P_{i,min}, P_{i,max}])$. Then the qMC experiment $\mathcal{E}(N, P)$ can be

Table 6.1: Distributions on the parameters for the qMC experiment.

Parameter	Distribution
l	unif($[0.05, 2]$)
ξ	0.1
h	0.02
$U_{19.5}$	unif($[10, 20]$)

written

$$\mathcal{L}(N,P) = \begin{pmatrix} P_{1,1} & P_{1,2} & P_{1,3} & \cdots & P_{1,P} \\ P_{2,1} & P_{2,2} & P_{2,3} & \cdots & P_{2,P} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ P_{N,1} & P_{N,2} & P_{N,3} & \cdots & P_{N,P} \end{pmatrix} \tag{6.12}$$

such that each row $\tau_n = \{P_{n,1}, P_{n,2}, P_{n,3}, \ldots, P_{n,P}\}$ of $\mathcal{L}(N,P)$ is a randomly drawn set of parameters called a trial. Then for all N trials τ_n, the system to be studied is examined for the parameter set τ. For each trial, a response variable r_n is measured, and the experiment $\mathcal{E}$ ultimately yields a set of trials and response variables $\{\mathcal{E}(N,P), \mathcal{R}\}$. The response surface can then be reconstructed as a function of the experimental parameters.

6.4.2 Sampling and reconstruction of the response surface

In the case of the horizontal pendulum system, the parameters of interest are the radius of gyration[2] l, normalized damping $\xi = \frac{b}{m}$, height of the arm above the platform h, and windspeed U, forming the parameter set $\{l, \xi, h, U\}$. To conduct the qMC experiment, equation (6.9), where ϕ and its derivatives are defined as in equation (6.8), is simulated using MATLAB's `ode45` numerical integrator for $T_s = 800\frac{2\pi}{\sqrt{g/l}}$ seconds for each trial. Table 6.1 provides information on the distribution of each parameter. The RMS velocity is then calculated and stored as the response variable r_n corresponding to the point τ_n.

[2] Note that l is not the same as the length of the pendulum arm.

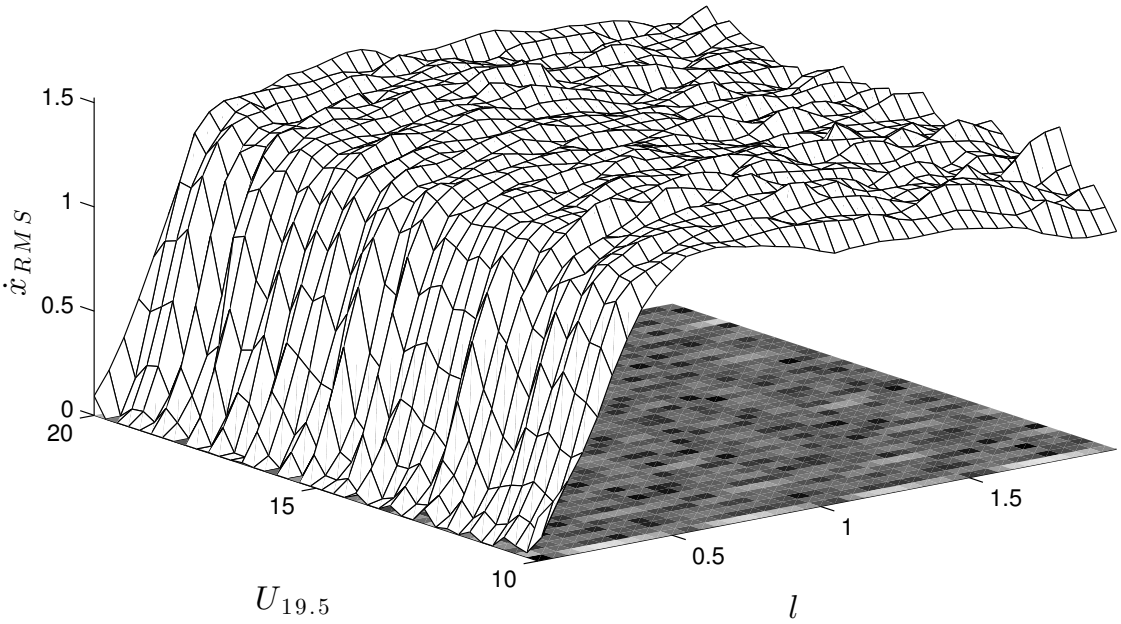

FIGURE 6.9: Response surface reconstruction by kriging. Units: $U_{19.5}$ (m/s), l (m), and $\dot{x}_{RMS}$ (m/s).

Given the three model parameters l (radius of gyration), $\xi = \frac{b}{m}$ (mass-normalized damping) and h (height of the arm above the platform), the one forcing parameter $U_{19.5}$, and the response data collected at each of these points, a response surface can be fit to the data and relevant statistics of the fit can be determined. Here, the response surface is reconstructed by kriging the data generated from the qMC experiment over a grid in the $\{l, U_{19.5}\}$ space, as can be seen in figure 6.9. This response surface can now be used to make informed design decisions for an experimental device intended for ocean testing.

6.4.3 Design recommendations

Figure 6.9 shows the estimated response surface kriged from data generated by the qMC experiment. From this information, and other observations made in the investigation of this device, a series of design recommendations can be made for a prototype design.

- The radius of gyration l should be chosen such that the design point is slightly greater than the peak value seen in figure 6.9 near $l = 0.8$. There is slight variation in the optimal value of l as wind speed varies, however, the most important consideration in choosing l is to ensure that the design point is not to the left of the peak value to account for any uncertainties that might drive the response to the lower shelf seen in the plot. Increasing l far beyond the peak value also sacrifices available RMS velocity and should be avoided.

- The normalized damping $\xi = \frac{b}{m}$ in a physical device would arise from both the mass of the rotating assembly and the characteristics of the rotating bearing. As observed in chapters 3 and 5, it is always beneficial to reduce damping if rotating behaviors are desired. This can be accomplished by adding mass to the rotating assembly if needed, but added mass would also affect the dynamics of the buoy in which the device was installed, potentially altering the response of the buoy in an unfavorable way or even sinking the entire assembly. Thus it is always advantageous to choose bearing with extremely low friction. Note that increasing ξ suppresses rotating behaviors and consequentially increases the optimal l. An exact relationship was not determined in this investigation.

- Adding electromechanical coupling will certainly have a drastic effect on the response of the system. Results shown in chapter 4 illustrate the effects of electrical loads on the frequency response of a nonlinear electromechanical system: the frequency at which peak power output occurs is shifted noticeably as the electrical load is varied. Again, the simplest strategy to mitigate these effects would likely be to add mass to the rotating assembly and increase the radius of gyration to provide additional torque on the generator; however, future work should investigate the effects of electromechanical coupling on the dynamics of the system.

- Increasing h, the height of the arm above the platform, typically seems to suppress the response of the system. Thus in designing a prototype the system should be mounted as near to the center of rotation as possible.

- Wind speed $U_{19.5}$ seems to have little effect on the response provided ξ is low and l is sufficiently large. Higher wind speeds do slightly increase the observed RMS velocity of the device, but the effect seems largely independent of l.

6.5 Conclusions

This chapter provided a model of random ocean waves and simulated the horizontal pendulum system in a stochastic environment. Analytical expressions for the wave angle and its derivatives were derived and used in the numerical investigations of the model. A quasi-Monte Carlo experiment was performed to characterize the response of the device over a typical range of wind speeds and a reasonable parameter range for a physical device. From these observations, preliminary design objectives are constructed to aid in the construction of a physical experiment.

7

Visualizing a Decision Boundary for Prediction of the Price Movements of a Security

As a result of the rapid growth of high-frequency trading in financial markets, characterization of market behaviors of very short time scales has become an area of great interest in quantitative finance [67]. One technique commonly used to study financial time series is regression modeling. Regression modeling is often employed in macroeconomic analyses to study the sensitivities of a macroeconomic variable of interest (e.g. the unemployment rate) to small changes in control parameters (e.g. interest rates). Hasbrouck and Saar brought the idea of vector autoregression to the analysis of financial time series in equities markets in reference [68], with the objective of determining the effects of small changes in the market structure on the price of a security.

Dr. Howle, David Martorana[1] and I collaborated to adapt the approaches presented in references [68, 69] to the development of a low-order discriminative classifier with the objective of using the classifier to determine a boundary in the market parameter space at which increasing price movements and decreasing price movements

[1] Duke B.S.E. in Mechanical Engineering and Economics, class of 2013

could be separated to a reasonable degree of accuracy. This resulted in the use of vector autoregression (VAR) and support vector machine (SVM) modeling and classification techniques to describe the evolution of a stock price on very short time scales (usually $t <= 1$ second) and to generate a set of classification variables from the variables with the highest information content. Leveraging the order management system built by Howle and documented in reference [70], we conducted an investigation of the 2011 trading year for the NASDAQ-tracking electronically traded fund (ETF), symbol QQQ. This chapter presents the results of our collaborative efforts in determining the information content of the limit order book and building a discriminative classifier in a subspace of the market parameters. In addition to verifying the results found in reference [69] and providing a discussion of the information content of the limit order book, this work contributes a visual representation of market behavior to a field that tends to eschew visual results in favor of tabulated numerical results. It is discovered that an effective ex-post-facto classifier[2] can be built in a limited two-dimensional subspace, and the law of supply and demand is empirically observed from the results.

7.1 Introduction

Over the past several decades there has been a marked increase in the complexity of financial markets. Sophisticated mathematical tools coupled with fast computers and the deluge of information available over the Internet have given rise to financial markets that evolve quickly in ways we struggle to comprehend. These changes have shaken the confidence of retail investors and made speculation even by institutional investors increasingly difficult and fiercely competitive [67]. Increasingly, algorithmic trading has come to dominate the financial markets in various forms ranging from

[2] The author cautions against the use of this method as trading strategy unless one desires to lose money as rapidly as possible.

ultra-low-latency arbitrage to sophisticated buying and selling algorithms used by large investors such as mutual funds to avoid price scalping [67].

In an effort to better understand the financial markets, traders have increasingly turned to statistical methods originating from the hard sciences in an effort to instill a sense of order in what is becoming an increasingly disordered business. This has given rise to the field of econophysics, which has its roots in Bachelier's study of cotton prices [71], the observations of Mandelbrot [72] regardning non-Gaussianity of price movements, options pricing (e.g. Black and Scholes [73]) and in statistical physics in general. The term itself was coined by Stanley in 1995 and the book by Mantegna and Stanley [74] is a comprehensive introduction to the topic. A book by Chakrabarti [75] gives a good overview of recent trends in the field.

The predecessor of econophysics is known as technical analysis. In the field of technical analysis [76], traders and brokers use historical stock market price, volume, news, and other information in an attempt to predict future price movements. Traditionally, technical analysis has relied on a mixture of mathematics, heuristics and luck to create profitable trading strategies. Of late, several technical indicators, including simple algorithms such as moving averages, volume weighted average prices, support and resistance levels or more complicated algorithms such as neural networks, support vector machines, or genetic algorithms are being studied in a more rigorous manner by econophysicists in an attempt to identify the buy and/or sell signals for a given financial instrument, and applied with great enthusiasm in the financial sector [67]. Indeed, high-frequency algorithmic trading (HFAT) has become more and more prevalent [77] with one estimate placing the volume of trades resulting from HFAT at 60% in the US equity markets for the 2009 market year [78].

However, despite advances in trading strategies, speed of execution, and the shift from broker-dealers to electronic matching engines, the fundamental principle of exchanges matching buyers and sellers of financial instruments has not changed [79].

In very general terms, stock orders may be described as either limit orders or market orders. For a limit order, the agent placing the order is guaranteed an execution price but is not guaranteed execution. In fact, for the data we analyzed in [70], for every order execution that occurred, there were 29.6 orders placed. For market orders, on the other hand, the agent placing the order has a guarantee of execution but no guarantee of order fill price [79].

The pool into which the orders are placed and through which the transaction occurs is frequently called an order book. There is a separate order book maintained for every financial instrument traded on a given stock exchange. Each order book can be further broken down into bid (buy) and ask (sell) books. A trade execution can occur either through a crossing limit order or through a market order. When a transaction occurs it is said to occur at the inside of the order book. That is, the order fills at the inside (most favorable to the agent placing the market or crossing order) price. These transactions at the inside price are the data frequently used in technical analysis. However, focusing only on the time and sales information, as is typically done with technical analysis, ignores the majority of the data available on the order book.

7.1.1 Organization of this investigation

The objective of this paper is to investigate the information content of order book information, including quote depth and order modifications, of the NASDAQ symbol QQQ, and to visualize a boundary between upticks and downticks in a subset of the the order book feature space. Our hope is that a visual representation will provide some insight into the market dynamics at a basic level. We approach this objective in two steps: first, we determine the information content of the features of the order book; then, we build a discriminative classifier by fitting a support vector machine to the data projected onto the two most informative features. This allows us to overlay

a decision boundary onto the projected data, hopefully yielding some intuition into the basic mechanisms of price formation.

Our first step is to determine the information content of the various features of the order book. We consider a wider range of features is typically present in the literature, to the best of our knowledge. Specifically, we look beyond transaction data and consider the information content of liquidity near the inside of the order book, as well as information regarding order modifications such as cancellations, deletions, replacements and trades. Several studies have investigated the information content of the limit order book of various securities with varying conclusions as to the information content of order book features. Some studies have focused on the statistical properties of the order book. Potters, Bouchaud and collaborators studied the statistical properties of the Paris Bourse and NASDAQ exchanges, including volume at depth and order cancellation [80, 81]. Challet and Stinchcombe studied the over-diffusive properties of prices on the NASDAQ Island order book [82]. Other studies have focused on regression modeling or other forecasting techniques to measure the information content of various market parameters. Kozhan and Salmon studied the economic value of exploiting order book information with a genetic algorithm trading strategy [83]. Cao et al. considered the information content of an open limit-order book and found that it had a moderate (22%) contribution to price discovery [84]. Hillman and Salmon considered a nonparametric approach in [85]. Harris and Panchapagesan found that the order book is informative in the decisions of NYSE specialists [86].

Of particular relevance to this investigation are the vector autoregressive models of Hasbrouck [68] and Mizrach [69]. An early investigation by Hasbrouck measured the information content of stock trades and determined the quote revision response to trades [68]. Mizrach expanded on the Hasbrouck model significantly by considering additional information present in the limit order book and determined the relevance

of order book information as well as the market impact of trades [69]. Our investigation seeks to expand on the model proposed by Mizrach, but rather than attempt to characterize the market impact of trades, we seek to determine if it is possible to form a set of features that can yield and effective 2D discriminative classifier for price movements. Specifically, we consider two vector autoregression (VAR) type models and draw conclusions regarding the significance of market data from the weightings found in each model. From our conclusions we design a set of two features and build a discriminative classifier for price movements using a support vector machine.

To that end, we have organized the paper as follows. First, the methods used to process and analyze the data are shown and discussed. We review a previously developed message processing system [70] that allows us to reconstruct the order book from NASDAQ TotalView-ITCH data. We then describe the Mizrach VAR model and our proposed principal component VAR model (PCVAR). We test each model on the entirety of the 2011 QQQ[3] dataset and generate weighting and significance estimates for each parameter. We observe similarities between the two models that indicate a parameter set which can be used to visualize a boundary between positive and negative price movements. Finally, we build a maximum-margin discriminative classifier in the subspace using a support vector machine (SVM) and interpret the results.

7.2 Data

Analyzing financial markets on short time scales is a big-data problem. Gigabytes of data are produced each day, and even the task of reassembling the time history of the order book requires an advanced software platform. In a previous investigation, we designed, programmed and evaluated an object-oriented order processing system built to parse NASDAQ TotalView-ITCH market data [70]. A brief review of the

[3] Note that for part of 2011 QQQ held the symbol QQQQ.

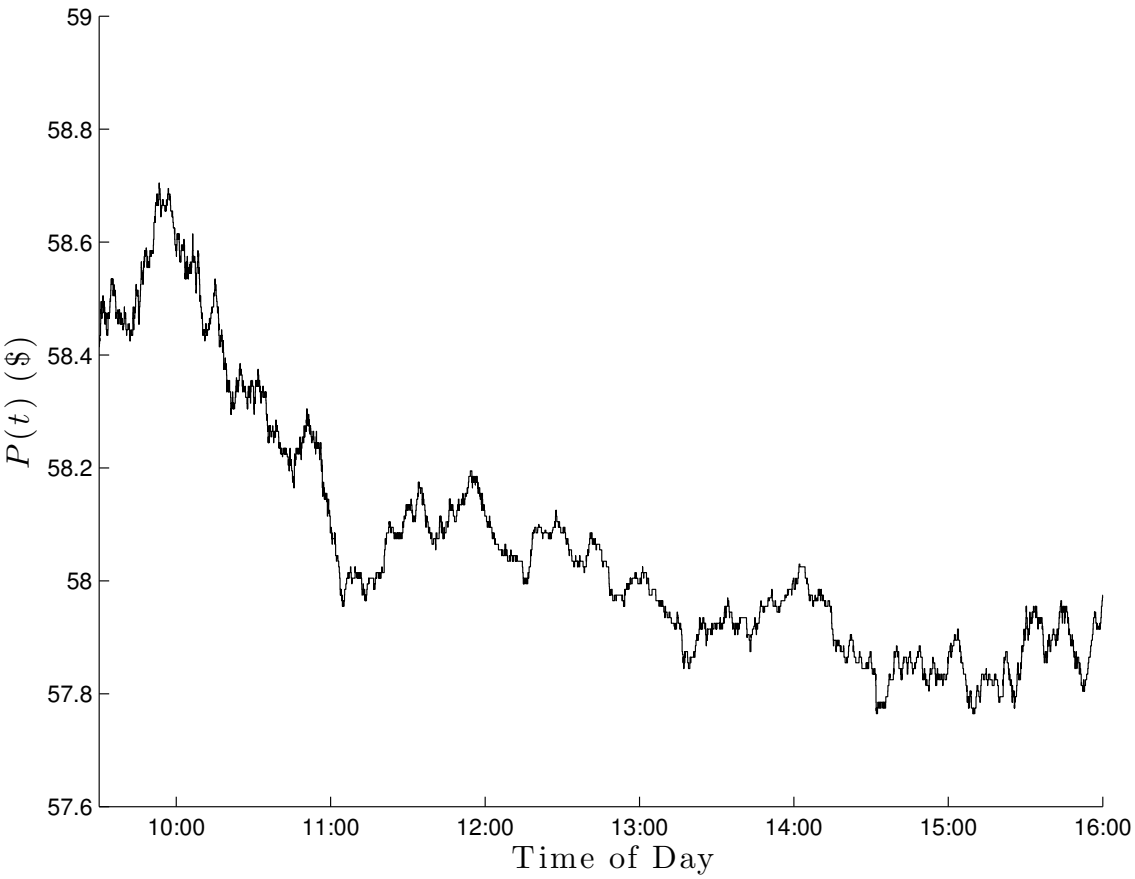

FIGURE 7.1: QQQ on 07 Jan 2011.

system is presented here for the convenience of the reader. The NASDAQ TotalView-ITCH system provides data in the form of messages; each message is encoded with identifying information including a type and any other relevant information pertaining to the message. A table of messages, their type, and daily frequency of appearance is shown in table 1. As can be seen from the table, an average of over 400 million messages per day are posted on by the NASDAQ TotalView-ITCH system. The sheer volume of messages has led many financial institutions to invest in high tech, bespoke hardware specifically tuned to the task of analyzing a particular exchange's data [87]. As we are not concerned with slight delays in data processing, and we lack the budget of a major financial institution, we sought to develop an object-oriented, easily reconfigurable system to parse the data and prepare it for statistical analysis.

Table 7.1: Message types (left column), average number of daily messages (center column) and standard deviation of the daily message count for the 251 trading days in the 2011 NASDAQ market year. The bolded message types affect the maintenance of the order book.

Message type	μ	σ
AddOrder	1.812e+08	6.641e+07
AddOrderMPID	7.989e+06	3.753e+06
BrokenTrade	9.976e+01	7.462e+02
CrossTrade	1.572e+04	9.008e+01
MarketParticipantPosition	1.822e+05	2.622e+04
NOII	1.132e+06	6.477e+03
OrderCancel	1.618e+06	5.748e+05
OrderDelete	1.829e+08	6.700e+07
OrderExecuted	7.523e+06	2.214e+06
OrderExecutedWithPrice	1.659e+05	1.053e+05
OrderReplace	3.849e+07	1.905e+07
RegSHORestriction	6.829e+03	2.916e+03
StockDirectory	7.996e+03	9.781e+02
StockTradingAction	8.005e+03	9.645e+02
SystemEvent	5.000e+00	0.000e+00
Timestamp	4.676e+04	7.152e+02
Trade	7.691e+05	2.205e+05
Total	4.220e+08	1.570e+08

For this task, we turned to the versatility of the Microsoft .NET platform, and developed an object-oriented order processing system in the C# language [88], which we will refer to as the NASDAQ Order Processing System (NOPS). NOPS is capable of either receiving messages in real-time or processing historical message databases. For the purposes of this investigation, we used our software solution to reconstruct the order book and stroboscopically sample the data in one second intervals. As messages arrive on a nanosecond timescale, this relatively long sample period greatly reduces the size of the data set, permitting easier analysis with user-friendly statistical tools such as those built into MATLAB [18].

FIGURE 7.2: Typical structure of the order book during the trading day. On the horizontal axis, orders are arranged by price. The vertical axis denotes the size of each order. Each order ranges in volume and its execution is prioritized by its time of arrival. The price level of an order is simply the distance of the order from the inside. There may be more than one order at a given price level. The central price is defined as the midpoint between the bid and ask sides of the book. NOPS is able to convert raw message data into market snapshots such as this.

7.2.1 Reconstructing the order book

As mentioned above, each message contains information regarding some aspect of the order book; by assembling and organizing the stream of messages, the order book can be constructed either in real time or from recorded message logs. After reconstructing the time history of the order book, the next step is to separate the orders into categories for analysis. The "inside" of the order book is where executions happen; this is what drives the price of security. On the inside, there are bid and ask orders. When these orders match, an execution occurs, and the price of security is updated. There are also orders away from the inside; in this investigation we consider not only the orders at the inside but orders up to four tiers away from the inside price.

Additionally, we examine the influence on price movement of order cancellations, deletions, replacements and executions. By post-processing and stroboscopically sampling the order book at one-second intervals, we can export a preliminary set of features that includes not only price and volume data but also order modification activity such as cancellations and replacements (c.f table 7.1).

7.2.2 Features of the order book and creation of feature matrices

The order book itself contains only a list of limit orders, with details for each order that include price, volume and priority data. We utilize NOPS to flatten this information into a list of features and response variables that can be used for modeling. We consider two response variables, r_t and x_t, the simple periodic return and the signed sum of executed trades, respectively. Let p_t^b be the inside bid price at time t and p_t^a be the inside ask price at at time t. Then we can define the simple periodic return as

$$r_t = \frac{1}{2}(p_t^b + p_t^a) - \frac{1}{2}(p_{t-1}^b + p_{t-1}^a), \tag{7.1}$$

i.e. the first difference of the central price. In addition to changes in price, it can also be useful to consider trade executions. Let x_t be the sum of signed trades executed between time t and $t-1$:

$$x_t = \text{Number of shares bought} - \text{Number of shares sold.} \tag{7.2}$$

We then define a series of features that may be used to forecast r_t and x_t. These features include quantity features, difference features, and other specifically defined features. Quantity or volume features are defined as

- $q_{k,t}^b$ ($q_{k,t}^a$), the quantity of shares at bid (ask) tier k at time t, with $k = 1$ at the inside;

- q_t^b (q_t^a), the total number of shares on the bid (ask) side of the order book;
- $q_t^{b,\mathrm{X}}$ ($q_t^{a,\mathrm{X}}$) the quantity of shares modified by the action "X", where "X" is add, delete, cancel, replace, trade, execute, or execute with price.

Difference features are constructed from quantity features as follows:

- $Q_t = q_t^b - q_t^a$,
- $Q_{k,t} = q_{k,t}^b - q_{k,t}^a$,
- $Q_t^x = q_t^{b,\mathrm{X}} - q_t^{a,\mathrm{X}}$.

We also introduce a demand function as defined in [69] as

$$D_t = \frac{\sum_{k=1}^{5} q_{k,t}^b \cdot p_{k,t}^b}{\sum_{k=1}^{5} q_{k,t}^b}; \tag{7.3}$$

which can be seen as a weighted average price on the bid side and is an indicator of demand (buying pressure).

We can then assemble a feature matrix with n of these features as follows:

$$\Phi = \begin{pmatrix} \phi_1(t_0) & \phi_2(t_0) & \phi_3(t_0) & \cdots & \phi_n(t_0) \\ \phi_1(t_1) & \phi_2(t_1) & \phi_3(t_1) & \cdots & \phi_n(t_1) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \phi_1(t_m) & \phi_2(t_m) & \phi_3(t_m) & \cdots & \phi_n(t_m) \end{pmatrix} \tag{7.4}$$

where $\Phi \in \mathbb{R}^{mxn}$, $t \in [t_0, t_m]$ and each $\phi_i(t)$ is a feature sampled at time t; each row is an observation in time. The choice of features is determined by the model, which in turn determines Φ.

7.3 VAR models

With data in hand, we now seek to formulate a model of the order book. Vector autoregression (VAR) models are a useful tool for understanding the information content of various aspects of market microstructure. Given a set of observations, a VAR(p) model forecasts future observations from linear combinations of the p most recent past observations and a set of weights fit to training data.

In [68], Hasbrouck considers several VAR models that relate quote revision r_t and trade executions x_t. In general, the models are of the form

$$r_{t+1} = a_{r,0} + \sum_{i=1}^{p} a_{r,i} r_{t-i} + \sum_{i=1}^{p} b_{r,i} x_{t-i} + \epsilon_{r,t} \tag{7.5a}$$

$$x_{t+1} = a_{x,0} + \sum_{i=1}^{p} a_{x,i} r_{t-i} + \sum_{i=1}^{p} b_{x,i} x_{t-i} + \epsilon_{x,t} \tag{7.5b}$$

where the $a_{(,i)}$ and $b_{(,i)}$ coefficients determine the impact of lagged information in predicting the next observation. In [69], Mizrach proposes an extension to the Hasbrouck/Saar VAR model

$$\begin{aligned} r_{t+1} = a_{r,0} &+ \sum_{i=1}^{5} a_{r,i} r_{t-i} + \sum_{i=1}^{15} b_{r,i} x_{t-i} + \sum_{k=1}^{5} \alpha_{r,k} \left(q^b_{k,t} - q^a_{k,t} \right) \\ &+ \sum_{i=1}^{15} \delta_{r,i} \Delta D_{t-i} + \gamma_r \left(q^b_t - q^a_t \right) \end{aligned} \tag{7.6a}$$

$$\begin{aligned} x_{t+1} = a_{x,0} &+ \sum_{i=1}^{5} a_{x,i} r_{t-i} + \sum_{i=1}^{15} b_{x,i} x_{t-i} + \sum_{k=1}^{5} \alpha_{x,k} \left(q^b_{k,t} - q^a_{k,t} \right) \\ &+ \sum_{i=1}^{15} \delta_{x,i} \Delta D_{t-i} + \gamma_x \left(q^b_t - q^a_t \right) \end{aligned} \tag{7.6b}$$

which makes use of the additional information present in the order book beyond the

autoregressive terms, including tiered volume differences $q^b_{k,t} - q^a_{k,t}$ and the demand function D_t.

7.3.1 PCVAR model

We seek to include the additional information at our disposal in a similar manner to the Mizrach and Hasbrouck/Saar VAR models; however we adopt a slightly different approach in doing so. Rather than include the features directly, we propose a dimensionality reduction to the first n principal components of the data. Given a matrix of observations Φ, we can compute its singular decomposition

$$\Phi = U\Sigma V^\top \tag{7.7}$$

from which we can determine the principal components (columns of V), the principal component scores $S = U\Sigma$, and the spectral power of each component $\Sigma^2_{i,i}$. We then form a VAR model from a subspace of the scores (i.e. the first n columns). As the principal components are eigenvectors of the covariance matrix $\Phi\Phi^\top$, they provide insight into which parameters are covariant. Creating a VAR model from the scores then allows us to determine the relative importance of each principal component in price formation. Then we seek a mixed VAR(p)/regression model of the form

$$r_{t+1} = a_{r,0} + \sum_{i=1}^{p} a_{r,i} r_{t-i} + \sum_{i=1}^{p} b_{r,i} x_{t-i} + \sum_{i=1}^{p} \left[\sum_{k=1}^{n} w^r_{k,i} s_{k,t-i} \right] + \epsilon_{r,t} \tag{7.8a}$$

$$x_{t+1} = a_{x,0} + \sum_{i=1}^{p} a_{x,i} r_{t-i} + \sum_{i=1}^{p} b_{x,i} x_{t-i} + \sum_{i=1}^{p} \left[\sum_{k=1}^{n} w^x_{k,i} s_{k,t-i} \right] + \epsilon_{x,t} \tag{7.8b}$$

where $s_k = \Phi \boldsymbol{c}_k$ and $\boldsymbol{c}_k$ is the k^{th} principal component.

7.3.2 General form of the models and fitting

Both the Mizrach and PCVAR models can be written as

$$Y_t = \sum_{i=1}^{p} \Phi_{t-i} A_i + \boldsymbol{\epsilon}, \tag{7.9}$$

or more compactly as

$$Y = Z\boldsymbol{w} + \boldsymbol{\epsilon}. \tag{7.10}$$

The matrix $Z \in \mathbb{R}^{m \times n \cdot p}$ contains n features and their p lags, sampled at m points in time t_0 through t_{m-1}, and the matrix $Y \in \mathbb{R}^{m \times r}$ contains the r response variables sampled at the m points in time t_1 through t_m. The weights, $\boldsymbol{w}$, are an $n \cdot p \times r$ matrix of coefficients corresponding to the n features with p lags and the r response variables in Y. The ordinary least-squares (OLS) estimate of the weights is given by

$$\boldsymbol{w} = (Z^\top Z)^{-1} Z^\top Y; \tag{7.11}$$

we use OLS to fit the weights for both models.

7.4 Results

For all 251 trading days in 2011, we applied both the Mizrach VAR model and the PCVAR model to arrive at weightings for each feature of the order book or each principal component. A report was generated containing a set of weights for each model on each day and summary statistics were generated for each model. Here we present both the summary statistics for the entire year and an example from January 2011.

7.4.1 Yearlong statistics

For the Mizrach VAR model, the results indicate that r_t is strongly anti-autocorrelated, x_t has a somewhat random effect on returns, D_t is strongly correlated with positive returns, and excess bid (ask) volume in the first few price tiers tends to drive the price up (down). Figure 7.3 gives a visual representation of these results and table 7.2 gives a table of the mean coefficients and their standard deviations.

For the PCVAR model, the results indicate that the first two to three principal components are the most informative, and that only one or two time lags are informative. We were able to compute an averaged representation of the principal components for the entire year, as shown in table 7.3. These components are interesting and yield some insight into the structure of correlations present in the market. From inspection of 7.3, one should notice a trend of nearly antisymmetric weightings across the order book. This indicates that the most important drivers of variance (and thus information content) are *differences* in information on the bid and ask sides of the order book. Considering table 7.4, the components are weighted such that the most significant drivers of price movements appear to be the differences in liquidity at various price tiers near the inside of the book, with a significant correlation time of only a few seconds. Note that the components c_1 through c_5 are ordered by their weights in table 7.3. Figure 7.4 gives a visual representation of each component's weight for up to five time lags.

We observe that demand-like variables that measure either a difference in liquidity on the bid and ask sides of the book or simply the liquidity available on the bid side tend to dominate both models. This is an interesting result as it indicates that the basic principle of supply and demand is tied to price formation (as one might readily expect), and that despite two different approaches to the problem the statistics indicate that demand-like variables are the most informative with respect

to changes in price. The weightings found in the Mizrach model are consistent with the findings in [69] and simply indicate that buying pressure drives price increases, and the weightings found in the PCAVAR model indicate that excess liquidity on the bid (ask) side drives price increases (decreases), all with $p < 0.01$.

Table 7.2: Summary of Mizrach VAR model parameters for 2011 QQQ data. The mean (standard deviation) for each coefficient is displayed with increasing time lags i on the horizontal. Results are truncated to the first five time lags.

τ_i	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$
$a_{r,i}(\times 1.0e4)$	-0.2206 (0.1572)	-0.1254 (0.0876)	-0.0803 (0.0539)	-0.0540 (0.0377)	-0.0323 (0.0244)
$b_{r,i}(\times 1.0e5)$	0.1028 (0.1189)	-0.0091 (0.0767)	-0.0111 (0.0760)	-0.0101 (0.0675)	-0.0084 (0.0702)
$\alpha_{r,1}(\times 1.0e4)$	0.1041 (0.0441)				
$\alpha_{r,2}(\times 1.0e4)$	0.0568 (0.0413)				
$\alpha_{r,3}(\times 1.0e4)$	0.0182 (0.0175)				
$\alpha_{r,4}(\times 1.0e4)$	0.0055 (0.0123)				
$\alpha_{r,5}(\times 1.0e4)$	0.0017 (0.0138)				
$\delta_{r,i}(\times 1.0e4)$	0.2234 (0.1659)	0.1262 (0.0863)	0.0777 (0.0479)	0.0533 (0.0344)	0.0319 (0.0245)

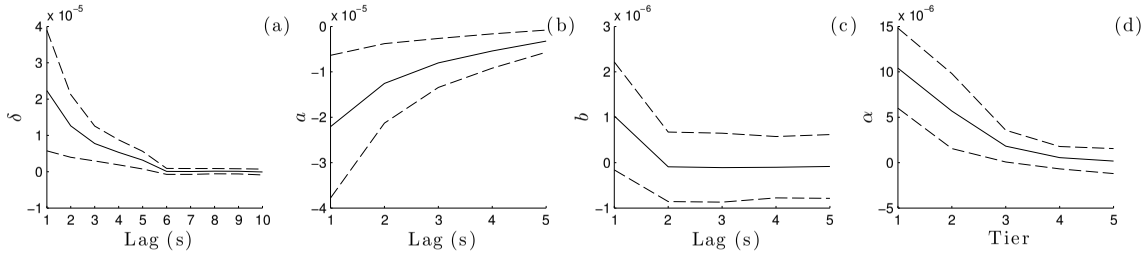

FIGURE 7.3: Mean and standard deviation influence of model parameters in the 2011 QQQ dataset. Shown: autoregression parameter for r_t (a), autoregression parameter for x_t (b), tiered, lagged volume differences $q^b_{k,t} - q^a_{k,t}$ (c), and the demand function D_t (d).

Table 7.3: Normalized principal components for the 2011 QQQ dataset, ranked in order of average weight by the PCVAR model. The first 95% of each component is explained by the terms in bold.

	$\boldsymbol{c}_1$		$\boldsymbol{c}_2$		$\boldsymbol{c}_3$		$\boldsymbol{c}_4$		$\boldsymbol{c}_5$	
	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask
q_t	**0.2811**	**-0.3382**	**0.6645**	**-0.7400**	**0.6886**	**-0.5777**	**0.8062**	**-0.1519**	**0.4414**	0.0258
$r_{1,t}$	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
$q_{1,t}$	**0.5322**	**-0.5530**	**0.0479**	**-0.0507**	**0.1301**	**-0.0699**	**0.2147**	**-0.0892**	**0.0859**	-0.0182
$r_{2,t}$	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
$q_{2,t}$	**0.3227**	**-0.3146**	**0.0195**	**-0.0223**	**0.0588**	**-0.0570**	**0.1208**	**-0.0560**	**0.0676**	**-0.0675**
$r_{3,t}$	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
$q_{3,t}$	-0.0331	0.0181	0.0043	-0.0079	0.0043	-0.0063	0.0335	0.0213	0.0116	0.0268
$r_{4,t}$	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
$q_{4,t}$	**-0.0378**	0.0330	0.0057	-0.0102	0.0043	0.0018	0.0228	0.0247	0.0093	0.0184
$r_{5,t}$	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000	-0.0000
$q_{5,t}$	-0.0259	0.0185	0.0090	-0.0130	0.0077	-0.0021	0.0205	0.0208	0.0115	0.0094
q_t^{added}	**-0.0414**	0.0299	-0.0142	**0.0284**	-0.0055	**0.1001**	-0.0047	**-0.0473**	**0.0703**	**0.6272**
$q_t^{\text{cancelled}}$	-0.0004	0.0000	0.0000	0.0001	0.0002	0.0004	0.0004	-0.0008	0.0050	0.0050
q_t^{deleted}	**-0.0601**	**0.0520**	-0.0106	**0.0222**	**0.0478**	**0.0493**	0.0132	**-0.0867**	**0.1283**	**0.5329**
q_t^{replaced}	0.0099	-0.0225	**0.0451**	**-0.0322**	**0.2069**	**-0.3253**	**0.4353**	**-0.2283**	**0.2501**	**-0.1565**
$\Delta q_t^{\text{replaced}}$	-0.0000	0.0004	0.0000	0.0001	-0.0036	0.0008	-0.0039	-0.0010	0.0001	-0.0023
q_t^{executed}	0.0015	-0.0017	-0.0001	0.0002	0.0039	-0.0007	0.0031	-0.0007	0.0161	0.0021
$q_t^{\text{ex. w/ p.}}$	-0.0000	-0.0000	0.0000	-0.0000	-0.0000	0.0002	0.0001	0.0004	-0.0003	-0.0000
q_t^{traded}	-0.0000	-0.0001	0.0000	0.0002	0.0000	0.0005	-0.0001	0.0003	0.0008	-0.0002

Table 7.4: Mean and standard deviation of the principal component weightings from the PCVAR model for the 2011 QQQ dataset.

τ_i	$i = 1$	$i = 2$	$i = 3$	$i = 4$	$i = 5$
$w_{1,i}(\times 1.0e9)$	0.3481 (0.1720)	0.1478 (0.0915)	0.0856 (0.0653)	0.0734 (0.0495)	0.0658 (0.0405)
$w_{2,i}(\times 1.0e9)$	0.1653 (0.0979)	0.0783 (0.0512)	0.0453 (0.0338)	0.0351 (0.0248)	0.0380 (0.0253)
$w_{3,i}(\times 1.0e10)$	0.7648 (0.6044)	0.3882 (0.2634)	0.2549 (0.1968)	0.2051 (0.1633)	0.2024 (0.1383)
$w_{4,i}(\times 1.0e10)$	0.3344 (0.3645)	0.1658 (0.1348)	0.1276 (0.1053)	0.1114 (0.1024)	0.1064 (0.0825)
$w_{5,i}(\times 1.0e10)$	0.1125 (0.0925)	0.0626 (0.0586)	0.0567 (0.0639)	0.0488 (0.0490)	0.0459 (0.0463)

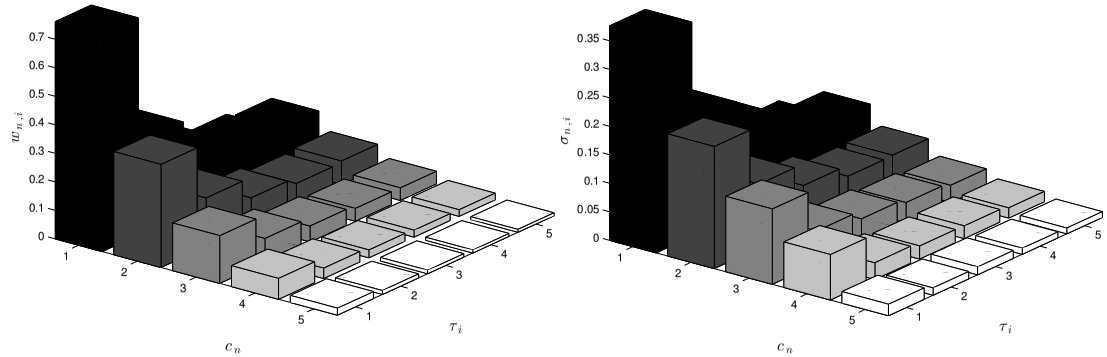

FIGURE 7.4: Left: Normalized mean yearly weights for the principal components defined in table 7.3. Right: Normalized standard deviation for the weights. Mean relative weights (left) and standard deviations (right) are shown against components (c_n axis) and time lags (τ_i axis).

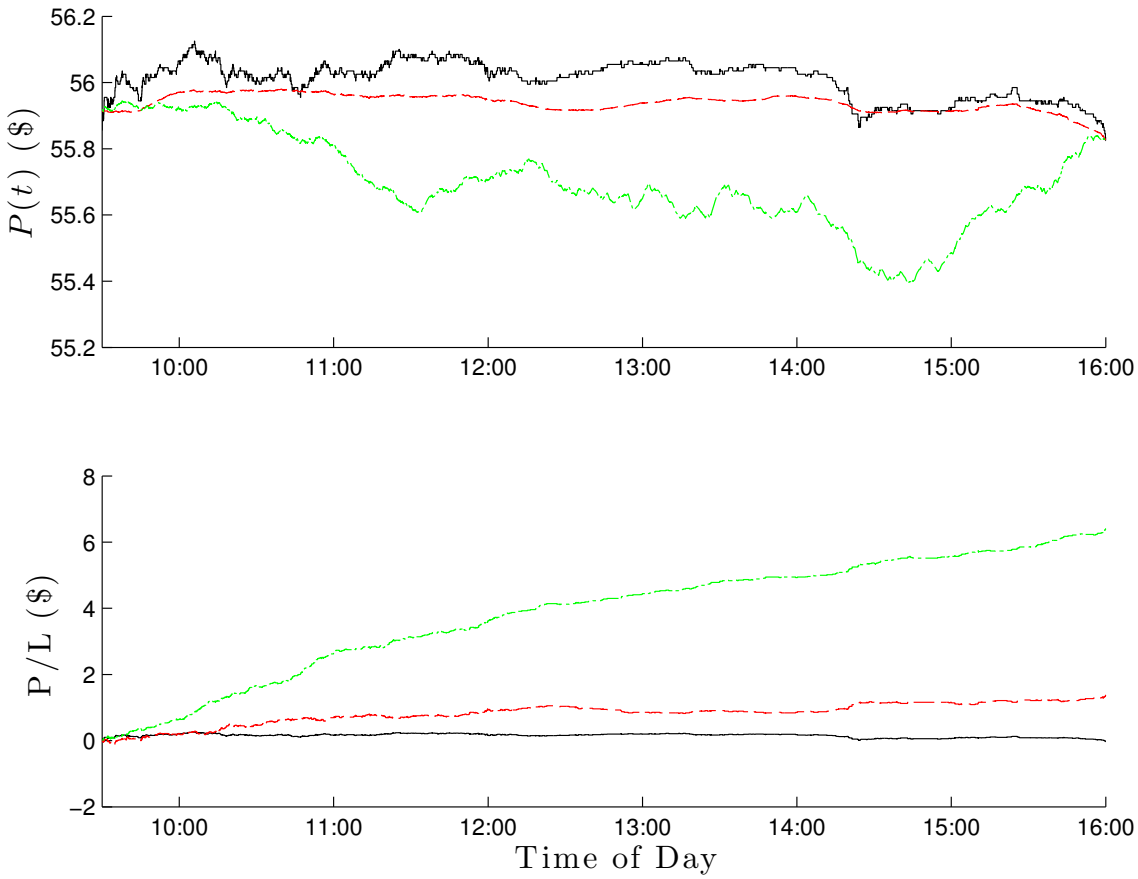

FIGURE 7.5: Example model outputs for 07 Jan 2011. The upper graph shows the price of the equity over a one-day period (solid black line), as well as in-sample predictions from the Mizrach VAR model (dashed red line) and the PCVAR model (dash-dotted green line). The lower graph shows the profit or loss of adhering to one of three strategies: buy and hold (solid black line), trade on the predicted $|x^0| \geqslant 0$ from the Mizrach VAR model (dashed red line), and trade on the predicted $|x^0| \geqslant 0$ from the PCVAR model (dash-dotted green line). Positions were held for one second and transaction costs were neglected.

7.4.2 Example: 07 January 2011

Given the fact that demand-like variables that measure the difference in liquidity tend to dominate in both the Mizrach and PCVAR models for the entire year's data set, we expect to find similar results when we consider a single trading day. To test this hypothesis, we considered all 251 trading days in 2011 and fit both the Mizrach and PCVAR models to every day, generating weights, p-values and figures for each

Table 7.5: Table of relative weights for the Mizrach VAR model, r_t. Weights are normalized to 1. Insignificant weights not shown. Time lags τ_i are shown as columns. See equation (7.6) for context.

τ_i	$i=1$	$i=2$	$i=3$	$i=4$	$i=5$
Bias	-	-	-	-	-
$a_{r,i}$	-1.00	-0.64	-0.48	-0.32	-0.16
$b_{r,i}$	-0.03	-	0.04	0.08	-
$\alpha_{r,0}$	0.57	-	-	-	-
$\alpha_{r,1}$	0.32	-	-	-	-
$\alpha_{r,2}$	0.15	-	-	-	-
$\alpha_{r,3}$	-0.02	-	-	-	-
$\alpha_{r,4}$	-0.07	-	-	-	-
$\delta_{r,i}$	1.00	0.66	0.47	0.28	0.23
γ_r	0.13	-	-	-	-

day. As the resulting report is too long to include here, we have included one day, 07 January 2011, as a sample to illustrate the approach.

Mizrach VAR Model

We begin by fitting the Mizrach VAR model. The parameters for r_t are shown in table 7.5. The most significant parameters are a and δ, indicating that returns are anti-autoregressive ($a < 0$) but strongly correlated with demand ($\delta > 0$). The tiered volume differences α_k were also significant, but primarily near the inside as the contribution from higher tiers drops off rapidly. We found the Spearman rank correlation coefficient to be $r = 0.174$ with $p < 0.01$.

PCVAR Model

Next we consider the PCVAR model. First, we determine the principal components $\boldsymbol{c}_n$, which are shown in table 7.6. We note that there is a high degree of symmetry in all of the principal components shown with respect to the bid and ask sides of the order book: each component seems to include a balance of effects from each side of the book. The first principal component represents the difference in total bid and

Table 7.6: Table of relative principal component weights for the first five components. Small components are indicated with a dash.

	$\boldsymbol{c}_1$		$\boldsymbol{c}_2$		$\boldsymbol{c}_3$		$\boldsymbol{c}_4$		$\boldsymbol{c}_5$	
	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask	Bid	Ask
q_t	-0.79	1.00	1.00	0.79	-	-	-	-	-	-
$r_{1,t}$	-	-	-	-	-	-	-	-	-	-
$q_{1,t}$	-	-	-	-	-	-	-0.85	1.00	-0.24	-
$r_{2,t}$	-	-	-	-	-	-	-	-	-	-
$q_{2,t}$	-	-	-	-	-0.14	-	-0.69	0.52	-	-0.20
$r_{3,t}$	-	-	-	-	-	-	-	-	-	-
$q_{3,t}$	-	-	-	-	-	-	0.16	-0.21	-0.13	-
$r_{4,t}$	-	-	-	-	-	-	-	-	-	-
$q_{4,t}$	-	-	-	-	-	-	-	-	-0.10	-
$r_{5,t}$	-	-	-	-	-	-	-	-	-	-
$q_{5,t}$	-	-	-	-	-	-	-	-	-	-
q_t^{added}	-	-	-	-	0.95	0.95	-	-	-0.92	1.00
$q_t^{\text{cancelled}}$	-	-	-	-	-	-	-	-	-	-
q_t^{deleted}	-	-	-	-	1.00	0.98	-	-0.12	-0.93	0.96
q_t^{replaced}	-	-	-	-	0.51	0.50	-	-	-0.34	-
q_t^{executed}	-	-	-	-	-	-	-	-	-	-
$q_t^{\text{ex. w/ p.}}$	-	-	-	-	-	-	-	-	-	-
q_t^{traded}	-	-	-	-	-	-	-	-	-	-

ask volume, the second represents the total volume open, the third represents total order modifications, the fourth represents tiered volume differences, and the fifth represents the difference in order modifications across the book.

It is crucial to note that the spectral power of each principal component is not necessarily related to its significance in the PCVAR model. The model indicates that the fourth component (tiered volume differences) is the most significant, followed by the first component (total volume difference). The second and third components (total volume and total order modifications, respectively) are insignificant, and the fifth component (total order modification difference) is somewhat significant. From this we conclude that volume imbalances, specifically near the inside, drive the price up (down) if there are more shares on the bid (ask) side of the book. Additionally, if there is a large turnover of shares on the bid (ask) side compared to the ask

Table 7.7: Table of relative weights for PCA regression, returns model. The notation $w_{n,i}$ means the weight of the n-th principal component at time lag i. Insignificant weights not shown; weights are normalized to 1.

	$w_{n,0}$	$w_{n,1}$	$w_{n,2}$	$w_{n,3}$	$w_{n,4}$
x_1	-0.40	-	-	-	0.14
x_2	-	-	-	-	-
x_3	-	-	-	-	-
x_4	-1.00	0.20	-	-	0.12
x_5	-0.11	-	-	-	-

(bid) side, the price tends to go up (down). For the PCVAR model, we found the Spearman rank-correlation coefficient to be $r = 0.306$ with $p < 0.01$. The results of each model and the P/L from a "trading strategy" that employed each model are shown in Fig. 7.5. Note that the trading strategy only held positions for one second, assumed no transaction costs, and did not account for its own impact on the market behavior, so it is not an accurate measure of what one might expect to gain or lose during a trading day using either of the VAR models as a forecasting tool; rather, it is an indicator of how informative the models are relative to the baseline of "buy-and-hold".

7.5 Determining a classification boundary

With the insights we gathered from the analysis of the Mizrach VAR model and the PCVAR model, we can see that the demand-like variables that display a difference in liquidity between the bid and ask sides of the book are the primary drivers of price movements. Hence it seems natural to construct artificial variables Q_k such that $Q_k = q_k^b - q_k^a$, i.e. the difference in bid and ask volume at the k^{th} tier. Our purpose in doing so is to (hopefully) find a space of two or three Q_k's that provide clear separation for uptick and downtick samples scattered in that space. We stress that this is an exercise in visualization, not forecasting, and the results we seek to find are a clear, visual boundary separating the two classes (upticks and downticks)

in a low-dimensional space rather than a quantitative representation of their division in an arbitrary space.

Considering Q_k for k ranging from 1 to 5, we found that Q_1 and Q_2 offered the greatest separation between observations of up and downticks. For $k \geqslant 3$, there was little visible separation. However, the separation shown in Q_1 and Q_2 was sufficient to build a discriminative classifier using a support vector machine with a radial basis function kernel. The results of the SVM test are shown in Fig. 7.6.

We found it particularly exciting that the decision boundary was simple. From Fig. 7.6, we can almost define a heuristic for predicting the direction of a tick based on the relative buy and sell volume at the first tier, i.e. the law of supply and demand. For $Q_1 > 0$, the next tick will likely be an uptick, and vice-versa. This is of course not a novel observation in economics, but it is exciting nonetheless to recover this observation from our data.

7.6 Conclusions

Our findings indicate that the primary driver of price movement is demand. Whether demand is modeled as a volume-weighted price or as a difference in liquidity seems not to make a difference: in either case, an imbalance in liquidity is indicative of an upcoming price movement. When there is excess liquidity on the buy (sell) side of the order book, the price tends to go up (down). This is an intuitive result that is consistent with the basic economic principle of supply and demand.

Principal component analysis of the market data sets yielded a surprisingly consistent set of principal components for all 251 trading days in 2011. Qualitatively, the first five principal components were almost exclusively (1) differences in liquidity at the first two to three price tiers, (2) total difference in liquidity on the bid and ask sides of the book, (3) differences in order modifications between the buy and sell sides of the book, (4) total trading volume, and (5) total order modifications.

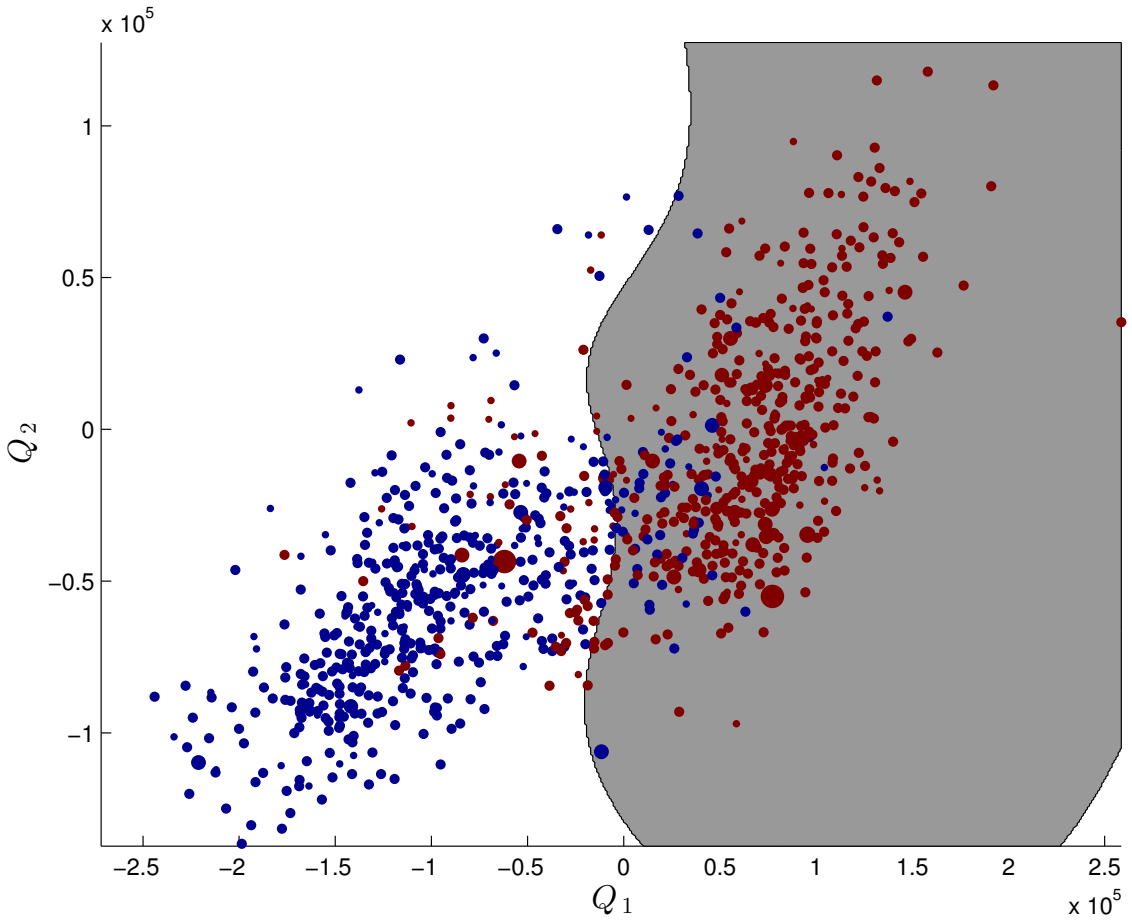

FIGURE 7.6: Decision boundary as determined by SVM. A coordinate system Q_1 and Q_2 was specified based on observations, with $Q_1 = q_1^b - q_1^a$ and $Q_2 = q_2^b - q_2^a$. In this space of near-inside demand-like variables, there is a clear separation of upticks (red) and downticks (blue) as denoted by the gray region (upticks) and the white region (downticks). The size of the dot indicates the volume of shares executed. The white area in the lower right hand corner is an artifact of the radial basis function kernel used in the SVM as there is no nearby data. Samples taken from 07 Jan 2011.

The PCVAR model consistently ranked the importance of the principal components in the order listed in the previous sentence. This is intriguing, as it seems to suggest that imbalances in information on the buy and sell sides of the order book are the most informative (more so than the amount of total information). However, we note that our analysis is not capable of determining whether this pattern drives price movements or is driven by price movements: we have only determined that this

information is correlated with future price movements.

We were able to use this information to generate two artificial features Q_1 and Q_2, which represent the difference in liquidity between the buy and sell sides of the order book at the inside and the first tier respectively. Utilizing these features we were able to create a scatterplot of the data with clear separation between upticks and downticks. We utilized a support vector machine with a radial basis function kernel to form a decision boundary that maximizes the separation between the two classes while minimizing misclassification rates. The boundary is simple and seems to indicate that Q_1, the difference in liquidity at the inside is a significantly more reliable classifier (less overlap) than Q_2, the difference in liquidity at the first tier. Plotting this decision boundary allows us to gain some visual intuition into the market structure of price movements. From this information we conclude that the simplest predictor of price movement is simply the sign of Q_1.

We note that our research does not attack the problem of when the next tick will occur, and that the "trading strategy" shown in Fig. 7.5 assumes that a price movement will occur every second. This is of course a false assumption, and would need to be tackled before attempting to implement any model shown in a trading strategy.

8

Conclusions

8.1 Summary and Conclusions

This dissertation provides advances in several topics related to nonlinear energy harvesting. An anharmonic averaging method for strongly nonlinear systems is developed and compared with the conventional harmonic averaging method. An efficient numerical approach to the study of the nonlinear energy harvester equations is designed and employed to study the nonlinear piezoelectric inertial generator. A new device for energy harvesting is proposed and studied to facilitate the development of an experimental prototype. A side investigation is conducted into the behavior of equities markets on short time scales, allowing for visualization of the market dynamics.

In chapter 2, a standardized modeling and analysis approach was presented that is used throughout the dissertation for all physical systems. Important theorems are shown and proved and a framework is provided for the discussions in chapters 3-5.

Chapter 3 arose from an attempt to apply averaging methods to the system in chapter 5. In so doing, results were found regarding the efficacy of harmonic and

anharmonic generating solutions in describing the behaviors of rotating systems. It is demonstrated that the choice of generating solution determines what results may be obtained: unless a generating solution that describes rotating behaviors is used, no information regarding rotating behaviors is obtained.

Chapter 4 applies numerical methods to the analysis of energy harvesting equations that have as yet only attracted analytical approaches in the literature. The chapter demonstrates that numerical methods can be applied in a systematic way to achieve the same results as analytical results have in the literature, but without the need for a possibly difficult-to-compute analytical solution.

Chapters 5 and 6 detail the motivation for and development of an ocean energy harvesting device. Characteristic behaviors of the device are shown for both harmonic and random forcing, and the investigation is focused on behaviors that would be observed should the device be placed in an ocean environment. A design study is conducted and a set of design recommendations is made.

Finally, chapter 7 investigates the effects of market microstructure on price formation. The chapter defines an effective classification boundary for price movements motivated both by observed market behaviors and the law of supply and demand.

8.2 Future Directions for Research

A multitude of future investigations could be conducted based on the research presented within this dissertation. Future contributions are broken down into four areas: (1) anharmonic averaging methods, (2) numerical continuation of periodic orbits in energy harvesting, (3) the horizontal pendulum system, and (4) research in financial markets.

8.2.1 Anharmonic averaging methods

Anharmonic averaging methods, although extremely challenging to implement, offer a rich description of behaviors in strongly nonlinear systems. Some future directions include:

1. Generalization of the method to systems of coupled oscillators. Mathematical details involving the convolution of Fourier series with mismatched periods will require effort to ensure usable results.

2. Generalization of the method to functions with arbitrary numbers of periods. Elliptic functions are biperiodic and are capable of describing systems with third-order polynomial restoring forces and trigonometric restoring torques. Multiperiodic (> 3) functions such as the hyperellpitic Abelian functions would be required to describe systems with more complex restoring forces or torques. Alternatively, a semi-numerical method could be devised that relied on Fourier expansion of the periodic, non-harmonic terms.

8.2.2 Numerical analysis of energy harvester models

As a whole, the energy harvesting literature tends to focus on analytical methods to find periodic orbits of energy harvesting systems and numerical methods to conduct phenomenological investigations of devices. As the complexity of models increases, analytical methods tend to become unwieldy, so a structured numerical approach is a viable alternative. Possible research opportunities include:

1. Investigation of the subharmonics of energy harvesting devices. Finding subharmonic orbits with analytical methods is a very involved process, but with a numerical approach finding subharmonics should be a much easier task. As some energy harvesting devices may be designed to harvest energy from a subharmonic, this research deserves further investigation.

2. Investigation of the response with multifrequency excitation. Energy harvesting devices often operate in environments with multiple driving frequencies. In nonlinear systems, this can drastically affect the response of the system. An analysis in a multifrequency excitation environment should be conducted in a future investigation.

3. Higher-fidelity uncertainty quantification using spectral methods. Finite difference approximations are a locally-valid approximation of the uncertainty. The assumption made is that the uncertainty varies linearly and is normally distributed; both assumptions are not necessarily the case in physical systems. For example, normally distributed errors in the length of a cantilever beam are nonsensical, as there is zero probability of an infinitely long or infinitely short beam. Spectral methods would allow for uniform distributions to be used to approximate the errors due to tolerances in the design of the device more accurately.

4. Uncertainty quantification specifically focused on physical parameters. In the investigation shown within the dissertation, uncertainties are taken with respect to the dimensionless parameters. In the development of an energy harvesting device, a more valid metric would be to compute the uncertainties in the physical parameters of the device such as magnet strength or component sizes.

5. Analysis of models incorporating more physically accurate magnet models (and thus more complex restoring forces) and more physically accurate damping models. Numerical methods permit higher-fidelity representations of nonlinear behaviors such as magnet-magnet interactions and structural damping effects that affect the performance of these devices in subtle but important ways.

8.2.3 Development of an ocean energy harvester

The primary focus of this research was to begin the development of an ocean energy harvester based on the horizontal pendulum system. As noted in the dissertation, analyzing this device proved to be an extremely challenging task. There is ample future work in the development of this device, including:

1. Further statistical investigations into the random vibrations of the device under Pierson-Moskowitz or similar wave forcing. As any device will ultimately operate in an ocean environment, it is logical to generate results with ocean-wave-like forcing. A large-scale numerical investigation on the Duke Scalable Computing Resource Center cluster is recommended.

2. Construction of a device and testing in the Dynamical Systems Laboratory wave flume. Preliminary results indicate that a sufficiently small device could be constructed and placed in the wave tank to experimentally verify the model behaviors. Work on the wireless data acquisition system begun in collaboration with the ME160 Blue Team and Marquese Pollard could be resumed to help in this investigation.

3. Development of a more detailed model, incorporating gyroscopic, pitch, heave, and buoyancy effects for a version of the device contained within a buoy. Analysis of this model will be computationally demanding and would best be completed in collaboration with the Dolbow laboratory.

8.2.4 Investigation of behaviors in financial markets

There is a nearly unending list of work to be done in characterizing the behavior of financial markets. Additional work should be performed in the following areas:

1. Characterization of the empirical distributions for all aspects of the limit order

book. Certain message types and behaviors may be more informative than others w.r.t. predicting price movements, so conditional distributions should be constructed and techniques from information theory should be used to analyze their predictive capacity.

2. Additional work in the visualization of market dynamics should be conducted. The visualization package developed by L. Howle should be reported in the literature alongside other results obtained over the course of this collaboration.

Appendix A

Properties of Elliptic Functions

In the course of searching for an approximate analytical solution to the horizontal pendulum system, efforts to describe the rotating solutions prompted an investigation of the utility of a higher-order averaging method using elliptic functions. The motivation for using elliptic functions as generating functions for an averaging transformation is their ability to describe rotating behavior, a high-energy behavior exhibited by the horizontal pendulum system that we believe to be advantageous for energy harvesting purposes. In order to more fully understand the approach, I conducted an investigation into the behavior of the librating and rotating solutions of the vertical pendulum under torque excitation using an approach described in [29, 27]. I found that a major barrier to understanding the approach lies in a lack of intuition with regards to the behaviors and graphs of the Jacobi elliptic functions. The objective of this appendix is to provide a mathematic and visual background for the Jacobi elliptic functions to aid the reader in understanding the approach taken in Chapter 3.

A.1 Jacobi Elliptic Functions

The Jacobi elliptic functions are a subset of elliptic functions, which are defined as the inverses of elliptic integrals (more on elliptic integrals to follow). The Jacobi elliptic functions are particularly useful because they are analogous to the trigonometric functions and can be used in the solution of several nonlinear ODEs.

A.1.1 Preliminary Definitions and Graphs

Jacobi elliptic functions arise from inverting the elliptic integral of the first kind[1],

$$u = \mathbf{F}(\phi, k) = \int_0^{\phi} \frac{\mathrm{d}t}{\sqrt{1 - k^2 \sin^2 t}} \tag{A.1}$$

where $k \in [0, 1]$ is the elliptic modulus, and ϕ is the Jacobi amplitude, which is actually more similar to an angular variable and is measured in radians. We define the Jacobi amplitude function as follows:

$$\phi = \mathbf{F}^{-1}(u, k) = \mathrm{am}(u, k). \tag{A.2}$$

In other words, $\mathrm{am}(u, k) = \mathrm{am}(\mathbf{F}(\phi, k), k) = \phi$: hence $\mathrm{am}(u, k)$ is the inverse of the integral. At $\phi = \frac{\pi}{2}$ the integrals are referred to as "complete": that is, $\mathbf{F}(\pi/2, k) \equiv \mathbf{K}(k)$, the complete elliptic integral of the first kind, and $\mathbf{E}(\pi/2, k) \equiv \mathbf{E}(k)$, the complete elliptic integral of the second kind.

Before continuing, it is important to obtain a visual understanding of these functions. Figure A.1 shows the complete elliptic integrals of the first and second kinds for values of $k \in [0, 1]$; figure A.2 shows surface plots of the incomplete elliptic integrals of the first and second kinds against their amplitude and parameter. More detailed

[1] Oddly enough it is the elliptic integral of the second kind, $\mathbf{E}(\phi, k) = \int_0^{\phi} \sqrt{1 - k^2 \sin^2 t}\, \mathrm{d}t$, that is used to calculate the arc length of an ellipse, but the elliptic integral of the first kind arises in a number of important contexts in physics.

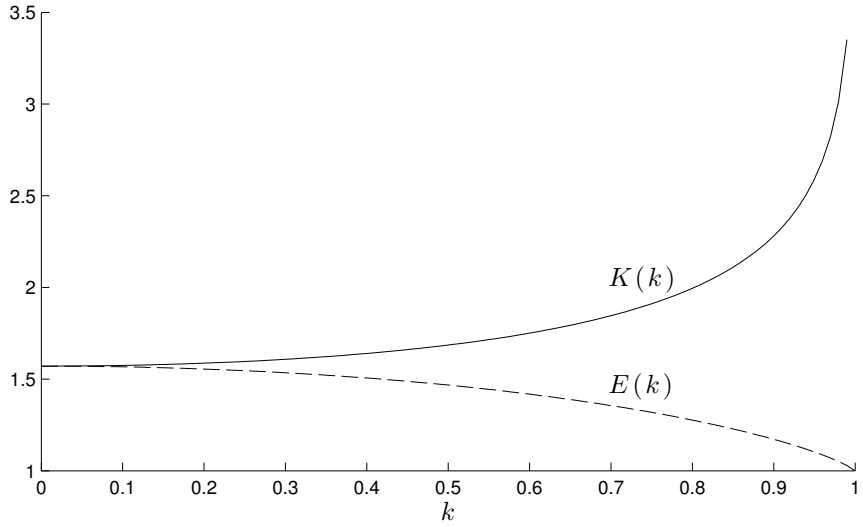

FIGURE A.1: Complete elliptic integrals of the first and second kind vs. the parameter k. In MATLAB, the calling sequence is `[K,E] = ellipke(m);` where $m = k^2$. In MAPLE, the calling sequence is `K = EllipticK(k);`, `E = EllipticE(k);`.

properties of elliptic integrals will be covered in later sections, but for the purpose of introducing the Jacobi elliptic functions this basic understanding will suffice.

A.1.2 The Jacobi Elliptic Functions and their Graphs

The Jacobi elliptic functions are all defined in terms of the Jacobi amplitude $\mathrm{am}(u,k)$. The function $\mathrm{am}(u,k)$ is periodic when restricted to the interval $[0, \mathbf{K}(k)]$, with period $2\,\mathbf{K}(k)$, and its behavior for various values of k can be seen in figure A.3 (a). In practice we do not treat the function as restricted to any interval as it is useful for modeling behaviors such as the angular position during continuous rotation of a pendulum.

The primary Jacobi elliptic functions are defined and read as follows:

$$\text{“sine amplitude”} \quad \sin \mathrm{am}(u,k) = \mathrm{sn}(u,k) \equiv \sin\phi \tag{A.3a}$$

$$\text{“cosine amplitude”} \quad \cos \mathrm{am}(u,k) = \mathrm{cn}(u,k) \equiv \cos\phi \tag{A.3b}$$

$$\text{“delta amplitude”} \quad \sqrt{1 - k^2 \sin^2(\mathrm{am}(u,k))} = \mathrm{dn}(u,k) \equiv \sqrt{1 - k^2 \sin^2\phi} \tag{A.3c}$$

(a) (b)

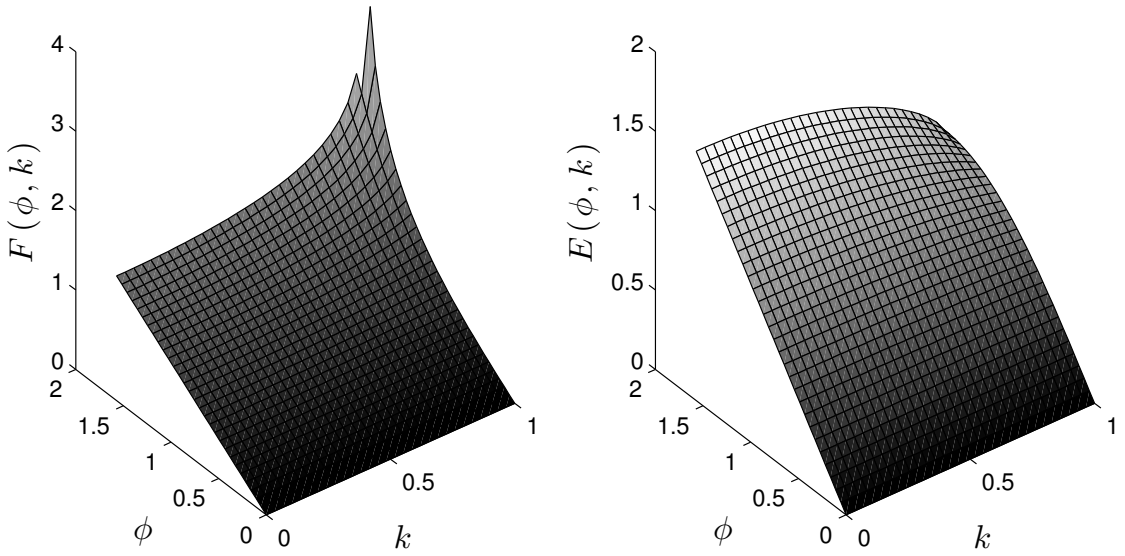

FIGURE A.2: Incomplete elliptic integrals of the first (a) and second (b) kind. In MATLABthe calling sequence is `[F,E] = elliptic12(`ϕ`,m)`, where $m = k^2$. This requires the elliptic package, see ref. [**?**]. In MAPLE the calling sequences are `F = EllipticF(`ϕ`,k);` and `E = EllipticE(`ϕ`,k);` respectively.

These functions are generalizations of trigonometric functions; to see this, set $k = 0$, which yields $u = \int_0^\phi \frac{\mathrm{d}t}{\sqrt{1-0}} = \phi$:

$$\mathrm{sn}(u, 0) = \sin u \tag{A.4a}$$

$$\mathrm{cn}(u, 0) = \cos u \tag{A.4b}$$

$$\mathrm{dn}(u, 0) = 1 \tag{A.4c}$$

Similarly, substituting $k = 1$ yields $u = \mathrm{arctanh}(\sin(\phi))$, which leads to

$$\mathrm{sn}(u, 1) = \tanh u \tag{A.5a}$$

$$\mathrm{cn}(u, 1) = \mathrm{sech}\, u \tag{A.5b}$$

$$\mathrm{dn}(u, 1) = \mathrm{sech}\, u, \tag{A.5c}$$

so the functions are also generalizations of the hyperbolic trigonometric functions.

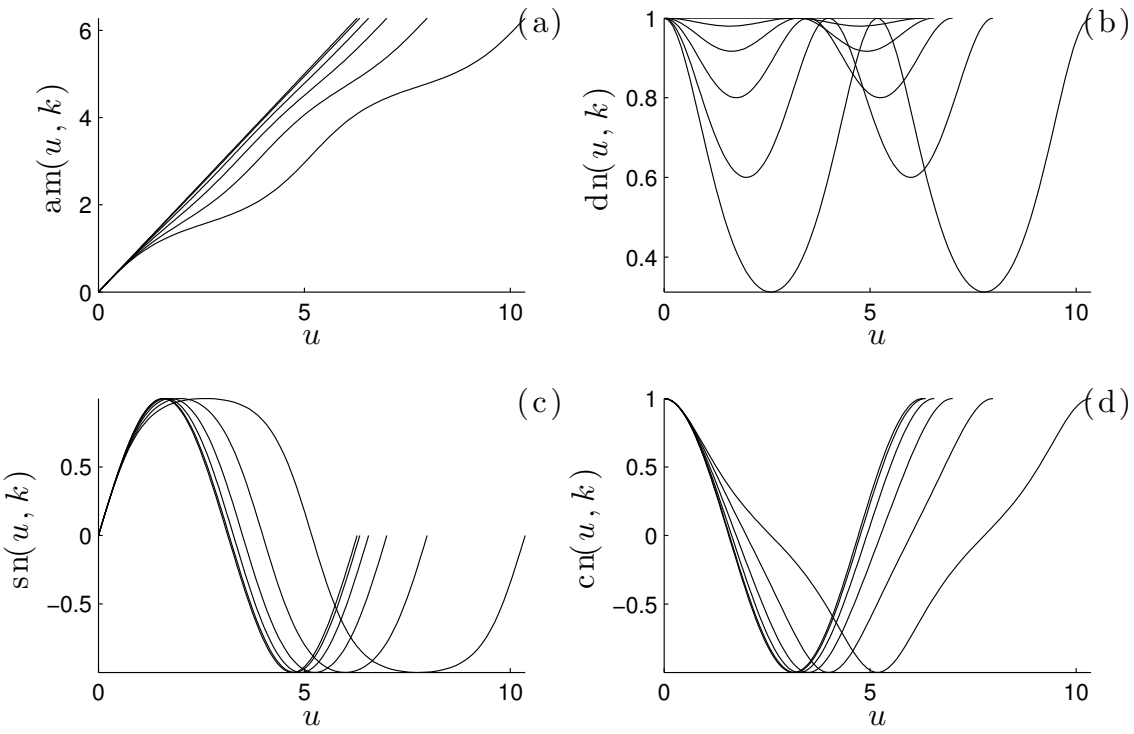

FIGURE A.3: The Jacobi amplitude (a), delta amplitude (b), sine amplitude (c) and cosine amplitude (d) functions for various values of k, with $k \in \{0, 0.2, 0.4, 0.6, 0.8, 0.95\}$ increasing as the function tends towards the right. Note that the periods $T = 2\,\mathbf{K}(k)$ (a,b) and $T = 4\,\mathbf{K}(k)$ (c,d) change as a function of k. Two periods of (a,b) are shown while only one period of (c,d) is shown.

Essentially, the Jacobi elliptic functions are a bridge between trigonometric and hyperbolic trigonometric functions. Plots of each function are shown in figure A.3.

Bibliography

[1] Willard J. Pierson Jr. and Lionel Moskowitz. A proposed spectral form for fully developed wind seas based on similarity theory of S.A. Kitaigorodskii. *Journal of Geophysical Research*, 69(24):5181–5190, 1964.

[2] B. Mann and N. Sims. Energy harvesting from the nonlinear oscillations of magnetic levitation. *Journal of Sound and Vibration*, 319:515–530, 2009.

[3] S. Stanton, C. McGehee, and B. Mann. Reversible hysteresis for broadband magnetopiezoelastic energy harvesting. *Applied Physics Letters*, 95(17):174103, 2009.

[4] S.C. Stanton, C.C. McGehee, and B.P. Mann. Nonlinear dynamics for broadband energy harvesting: investigation of a bistable piezoelectric inertial generator. *Physica D*, 239:640–653, 2010.

[5] S.C. Stanton, A. Erturk, B.P. Mann, and D.J. Inman. Resonant manifestation of intrinsic nonlinearities within electroelastic micropower generators. *Applied Physics Letters*, 97(25):254101, 2010.

[6] S. Bravo Yuste. Comments on the method of harmonic balance in which Jacobi elliptic functions are used. *Journal of Sound and Vibration*, 145(3):381–390, 1991.

[7] B.P. Mann, D.A.W. Barton, and B.A.M. Owens. Uncertainty in performance for linear and nonlinear energy harvesting strategies. *Journal of Intelligent Material Systems and Structures*, 0(0):1–10, 2012.

[8] S. Stanton, B. Owens, and B. Mann. Harmonic balance analysis of the bistable piezoelectric inertial generator. *Journal of Sound and Vibration*, 331:3617–3627, 2012.

[9] Jean-Luc Sabrier. *The Self-Winding Watch: 18th - 21st Century*. Editions Cercle D'art (Acc), 2012.

[10] Christopher M. Bishop. *Machine Learning and Pattern Recognition*. Springer Science + Business Media, 2006.

[11] Mark H. Holmes. *Introduction to the Foundations of Applied Mathematics*, volume 56 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 2009.

[12] Mark H. Holmes. *Introduction to Perturbation Methods*, volume 20 of *Texts in Applied Mathematics*. Springer-Verlag, 1995.

[13] Carl M. Bender and Steven A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers: Asymptotic Methods and Perturbation Theory*. Springer-Verlag, New York, 1999.

[14] E.J. Hinch. *Perturbation Methods*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Oxford, United Kingdom, 1991.

[15] Donald T. Greenwood. *Advanced Dynamics*. Cambridge University Press, Oxford, United Kingdom, 2006.

[16] I.M. Gelfand and S.V. Fomin. *Calculus of Variations*. Dover Publications, Englewood Cliffs, NJ, 1991.

[17] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes*. Cambridge University Press, New York, 3 edition, 2007.

[18] The MathWorks, Inc. MATLAB R2011b, 2011.

[19] D. Jordan and P. Smith. *Nonlinear Ordinary Differential Equations: An Introduction for Scientists and Engineers*. Oxford University Press, Oxford, United Kingdom, 2007.

[20] John Guckenheimer and Philip Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, volume 42 of *Applied Mathematical Sciences*. Springer-Verlag, New York.

[21] Jan A. Sanders, Ferdinand Verhulst, and James Murdock. *Averaging Methods in Nonlinear Dynamical Systems*, volume 59 of *Applied Mathematical Sciences*. Springer, 2nd edition, 2007.

[22] Ferdinand Verhulst. *Methods and Applications of Singular Perturbations: Boundary Layers and Multiple Timescale Dynamics*, volume 50 of *Texts in Applied Mathematics*. Springer-Verlag, New York, 2000.

[23] Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*, volume 112 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 3 edition, 2004.

[24] Rüdiger Seydel. *Practical Bifurcation and Stability Analysis*, volume 5 of *Interdisciplinary Applied Mathematics*. Springer Science + Business Media, New York, 3rd edition, 2010.

[25] Willy J.F. Govaerts. *Numerical Methods for Bifurcations of Dynamical Equilibria*. Society for Industrial and Applied Mathematics, Philadelphia, 1987.

[26] N. Kryloff and N. Bogoliuboff. *Introduction to Non-Linear Mechanics*. Princeton University Press, 1947.

[27] Ferdinand F. Cap. Averaging method for the solution of non-linear differential equations with periodic non-harmonic solutions. *International Journal of Non-Linear Mechanics*, 9:441–450, 1974.

[28] G. Pocobelli. Electron motion in a slowly varying wave. *Physics of Fluids*, 24(12):2173–2176, 1981.

[29] V. Roy. Averaging method for strongly non-linear oscillators with periodic excitations. *International Journal of Non-Linear Mechanics*, 29(5):737–753, 1994.

[30] T. Okabe and T. Kondou. Improvement to the averaging method using the Jacobian elliptic function. *Journal of Sound and Vibration*, 320:339–364, 2009.

[31] T. Okabe, T. Kondou, and J. Ohnishi. Elliptic averaging methods using the sum of jacobian elliptic delta and zeta functions as the generating solution. *International Journal of Non-Linear Mechanics*, 46:159–169, 2011.

[32] J. Garcia-Margallo and J. Diaz Bejarano. A generalization of the method of harmonic balance. *Journal of Sound and Vibration*, 116(3):591–595, 1987.

[33] S. Bravo Yuste and J. Diaz Bejarano. Amplitude decay of damped non-linear oscillators studied with Jacobian elliptic functions. *Journal of Sound and Vibration*, 114(1):33–44, 1986.

[34] S. Bravo Yuste and J. Diaz Bejarano. Improvement of a Krylov-Bogoliubov method that uses Jacobi elliptic functions. *Journal of Sound and Vibration*, 139(1):151–163, 1990.

[35] V. Coppola and R. Rand. *Macsyma program to implement averaging using elliptic functions*. Springer-Verlag, 1991.

[36] F. Lakrad and M. Belhaq. Periodic solutions of strongly non-linear oscillators by the multiple scales method. *Journal of Sound and Vibration*, 258(4):677–700, 2002.

[37] Paul F. Byrd and Morris D. Friedman. *Handbook of Elliptic Integrals for Engineers and Physicists*. Springer, 1954.

[38] Derek F. Lawden. *Elliptic Functions and Applications*, volume 80 of *Applied Mathematical Sciences*. Springer-Verlag, New York, NY, 1980.

[39] Alain J. Brizard. Jacobi zeta function and action-angle coordinates for the pendulum. *Communications in Nonlinear Science and Numerical Simulation*, 18:511–518, 2013.

[40] Maple 17. Maplesoft, a division of Waterloo Maple Inc., Waterloo, Ontario.

[41] R.L. Harne and K.W. Wang. A review of the recent research on vibration energy harvesting via bistable systems. *Smart Materials and Structures*, 22, 2013.

[42] S. Roundy, P. Wright, and J. Rabaey. A study of low level vibrations as a power source for wireless sensor nodes. *Computational Communications*, 26:1131–44, 2003.

[43] D.A.W. Barton, S.G. Burrow, and L.R. Clare. Energy harvesting from vibrations with a nonlinear oscillator. In *Proceedings of the International Design Engineering Technical Conference*, San Diego, 2009. ASME.

[44] A. Erturk, J. Hoffmann, and D. Inman. A piezomagnetoelastic structure for broadband vibration energy harvesting. *Applied Physics Letters*, 94(25):254102–254103, 2009.

[45] A. Erturk and D. Inman. Broadband piezoelectric power generation on high-energy orbits of the bistable duffing oscillator with electromechanical coupling. *Journal of Sound and Vibration*, 330(10):2339–2353, 2011.

[46] S. Stanton, A. Erturk, B. Mann, and D. Inman. Nonlinear nonconservative behavior and modeling of piezoelectric energy harvesters including proof mass effects. *Journal of Intelligent Material Systems and Structures*, 23(2):183–199, 2012.

[47] S. Stanton, A. Erturk, B. Mann, and D. Inman. Nonlinear piezoelectricity in electroelastic energy harvesters: modeling and experimental identification. *Journal of Applied Physics*, 108(7):074903–9, 2010.

[48] M.A. Karami and D.J. Inman. Equivalent damping and frequency change for linear and nonlinear hybrid vibrational energy harvesting systems. *Journal of Sound and Vibration*, 330:5583–5597, 2011.

[49] Lloyd N. Trefethen. *Spectral Methods in MATLAB*, volume 10 of *Software, Environments and Tools*. Society for Industrial and Applied Mathematics, Philadelphia, 2000.

[50] John P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover Publications, Mineaola, NY, 2 edition, 2001.

[51] Hugh W. Coleman and W. Gleen Steele. *Experimentation and Uncertainty Analysis for Engineers*. Wiley-Interscience, New York, 2 edition, 1999.

[52] O.P. Le Maître and O.M. Knio. *Spectral Methods for Uncertainty Quantification*. Scientific Computation. Springer Science + Business Media, New York, 2010.

[53] M.E. Henderson. Multiple parameter continuation: Computing implicitly defined k-manifolds. *International Journal of Bifurcation and Chaos*, 12(3):451–476, 2002.

[54] H. Cavendish. Experiments to determine the density of the earth. *Philosophical Transactions of the Royal Society of London*, 88:469–526, 1798.

[55] R. Peters. Mechanically adjustable balance and sensitive tilt meter. *Meas. Sci. Technol.*, 1:1131–1135, 1990.

[56] V. Graizer. Inertial seismometry methods. *Izvestiya, Earth Physics*, 27(1):51–61, 1991.

[57] V. Graizer. Effect of tilt on strong motion data processing. *Soil Dynamics and Earthquake Engineering*, 25:197–204, 2005.

[58] J. Hubbard. The forced damped pendulum: Chaos, complication and control. *The American Mathematical Monthly*, 106(8):741–758, 1999.

[59] M. Bartuccelli, G. Gentile, and K. Georgiou. On the dynamics of a vertically driven damped planar pendulum. *Proceedings of the Royal Society of London A*, 47:3007–3022, 2001.

[60] J. Scmitt and P. Bayly. Bifurcations in the mean angle of a horizontally shaken pendulum: Analysis and experiment. *Nonlinear Dynamics*, 15:1–14, 1999.

[61] S. Bishop, D. Xu, and M. Clifford. Flexible control of the parametrically excited pendulum. *Proceedings: Mathematical, Physical and Engineering Sciences*, 452(1951):1789–1806, 1996.

[62] E. Butcher and S. Sinha. Symbolic computation of secondary bifurcations in a parametrically excited simple pendulum. *Int. J. Bifurcation Chaos*, 8:627–637, 1998.

[63] M. Clifford and S. Bishop. inverted oscillations of a driven pendulum. *Proceedings of the Royal Society of London A*, 454:2811–2817, 1997.

[64] B. Mann and M. Koplow. Symmetry breaking bifurcations of a parametrically excited pendulum. *Nonlinear Dynamics*, 46:427–437, 2006.

[65] I.M. Sobol. Distribution of points in a cube and approximate evaluation of integrals. *Comput. Maths. Math. Phys.*, 7:86–112, 1967.

[66] Robert H. Stewart. *Introduction to Physical Oceanography*. Department of Oceanography, Texas A&M University, 2008.

[67] S. Patterson. *Dark Pools*. Crown Business, New York, 2012.

[68] J. Hasbrouck and G. Saar. Measuring the information content of stock trades. *Journal of Financial Markets*, 46:179–207, 1991.

[69] Bruce Mizrach. The next tick on NASDAQ. *Quantitative Finance*, 8(1):19–40, 2008.

[70] L.E. Howle, C.C. McGehee, and B.P. Mann. An object-oriented library for real-time processing of NASDAQ order book data. *JCEIT*, 1(2), 2012.

[71] L. Bachelier. Théorie de la spéculation. *Annales Scientifiques de l'École Normale Supèrieure*, 3(17):21–86, 1900.

[72] B. Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.

[73] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

[74] R.N. Mantegna and H.E. Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press, 2000.

[75] B. Chakrabarti. *Econophysics and sociophysics: trends and perspectives*. Weinheim, 2006.

[76] J.J. Murphy. *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York Institute of Finance, New York, 1999.

[77] G. Nuti, M. Mirghaemi, P. Traleaven, and C. Yingsaerre. Algorithmic trading. *Computer*, 44(11):61–69, 2011.

[78] R. K. Narang. *Inside the Black Box: The Simple Truth About Quantitative Trading*. Wiley Finance, Wiley, Hoboken, N.J., 2009.

[79] R.J. Teweles and E.S. Bradley. *The Stock Market*. Wiley, New York, 1998.

[80] M. Potters and J.P. Bouchaud. More statistical properties of order books and price impact. *Physica A*, 324:133–140, 2003.

[81] J.P. Bouchaud, Marc Mézard, and Marc Potters. Statistical properties of of stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256, 2002.

[82] D. Challet and R. Stinchombe. Limit order market analysis and modelling: On a universal cause for over-diffusive prices. *Physica A*, 324:141–145, 2003.

[83] R. Kozhan and M. Salmon. The information content of the limit order book: The case of an Fx market. *Journal of Financial Markets*, 15:1–28, 2012.

[84] Charles Cao, Oliver Hansch, and X. Wang. The information content of an open limit-order book. *The Journal of Futures Markets*, 29(1):16–41, 2009.

[85] R. Hillman and M. Salmon. Intrinsic stationarity: investigating predictability in real-time forex transactions. *Journal of Financial Forecasting*, 1:3–43, 2007.

[86] Lawrence E. Harris and Venkatesh Panchapagesan. The information content of the limit order book: evidence from NYSE specialist trading decisions. *Journal of Financial Markets*, 8:25–67, 2005.

[87] M. Millar. Lightning fast future traders working in nanoseconds. *BBC Mobile*, 2011.

[88] Microsoft Corporation. .NET 4.0 Framework, 2010.

Biography

Clark McGehee was born on July 15, 1988 in Atlanta, Georgia. He received his B.S.E. degree in Mechanical Engineering from Duke University in Durham, NC in 2010 and earned an M.Sc. in Mechanical Engineering from Duke University in 2011. In May 2010, he was commissioned as a second lieutenant in the United States Air Force and has been in reserve status during the completion of his doctoral program. He began collaboration with Dr. Brian Mann in the Dynamical Systems Laboratory as an undergraduate student in January 2009 as part of the Pratt Fellows program, led by Dean Martha Absher in the Pratt School of Engineering. Following collaboration with Dr. Mann and Dr. Samuel C. Stanton, he decided to remain at Duke and pursue a doctoral degree under the advising of Dr. Mann. Clark has participated in the publication of four journal articles as a second author and has published two first-author conference articles since beginning his graduate studies at Duke. He currently has three first-author publications in submission.

Following his PhD, Clark will enter active duty in the Air Force as an officer and developmental engineer in the Air Force Materiel Command.