ABSTRACT

MARTEN, ALEX LENNART. Essays on the Application and Computation of Real Options. (Under the direction of Professor Paul L. Fackler.)

This dissertation presents a series of three essays that examine applications and computational issues associated with the use of stochastic optimal control modeling in the field of economics. In the first essay we examine the problem of valuing brownfield remediation and redevelopment projects amid regulatory and market uncertainty. A real options framework is developed to model the dynamic behavior of developers working with environmentally contaminated land in an investment environment with stochastic real estate prices and an uncertain entitlement process. In a case study of an actual brownfield regeneration project we examine the impact of entitlement risk on the value of the site and optimal developer behavior. The second essay presents a numerical method for solving optimal switching models combined with a stochastic control. For this class of hybrid control problems the value function and the optimal control policy are the solution to a Hamilton-Jacobi-Bellman quasi-variational inequality. We present a technique whereby approximating the value function using projection methods the Hamilton-Jacobi-Bellman quasi-variational inequality may be recast as extended vertical non-linear complementarity problem that may be solved using Newton's method. In the third essay we present a new method for estimating the parameters of stochastic differential equations using low observation frequency data. The technique utilizes a quasi-maximum likelihood framework with the assumption of a Gaussian conditional transition density for the process. In order to reduce the error associated with the normality assumption sub-intervals are incorporated and integrated out using the Chapman-Kolmogorov equation and multi-dimensional Gauss Hermite quadrature. Further improvements are made through the use of Richardson extrapolation and higher order approximations for the conditional mean and variance of the process, resulting in an algorithm that may easily produce third and fourth order approximations for the conditional transition density.

© Copyright 2009 by Alex Lennart Marten

All Rights Reserved

Essays on the Application and Computation of Real Options

by
Alex Lennart Marten

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Economics

Raleigh, North Carolina

2009

APPROVED BY:

Dr. John J. Seater

Dr. Denis Pelletier

Dr. Paul L. Fackler
Chair of Advisory Committee

Dr. Roger H. von Haefen

DEDICATION

To my parents

BIOGRAPHY

Alex Marten was born in 1981 in Allentown, Pennsylvania. He is the son of Lennart and Lisa Marten and the brother of Britt Marten. Alex received his Bachelor of Science degree in Economics and Political Science from the University of Pittsburgh in 2003. He then moved to North Carolina to pursue a doctorate in Economics at North Carolina State University.

ACKNOWLEDGMENTS

First and foremost, I must thank my loving wife Christin for all of her encouragement and patience during this endeavor. Without her support this would not have been possible. I am also indebted to my advisor Paul Fackler for his guidance and insights. His direction and willingness to share his time and wisdom have been crucial to this effort. I would also like to thank the members of my committee, John Seater, Roger von Haefen, and Denis Pelletier for their thoughts and suggestions. Finally I would like to acknowledge the support and advice I have received from my friends and colleagues. Thanks to Brad, Chris, Erin, Jake, Metin, Robert, Ryan, and Tim.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

This dissertation presents a series of three essays that examine applications and computational issues associated with the use of stochastic optimal control modeling in the field of economics. In Chapter 2 we consider the problem of valuing brownfield remediation and redevelopment projects amid regulatory and market uncertainty. The question regarding the worth of an environmentally contaminated site has been the subject of great interest since the passage of the Comprehensive Environmental Response, Compensation, and Liability Act in 1980. This interest comes from private developers trying to select profitable investments in addition to those parties interested in designing effective incentives to encourage brownfield reuse. In either case a key component of understanding the value of such redevelopment projects is the ability to correctly account for the dynamic behavior of developers in uncertain investment environments. Previous efforts to model this scenario have made steady improvements, however they have consistently over simplified the problem and omitted important characteristics. We present a significant contribution to this area by working in consultation with a leading remediation firm to develop a more acurate real options framework for the value of brownfield remediation and redevelopment projects.

In order to achieve a realistic depiction of such investments we provide a number of extensions to the traditional approach to brownfield valuation. Instead of adhering to the typical simplifying assumption that prices follow geometric Brownian motion, we allow the temporal dynamics of real estate markets to exhibit short-term variation in the growth rate. This allows the framework to capture the effects of common market phenomenon often referred to as housing "bubbles" and "slumps". In addition, the framework includes the important but repeatedly overlooked feature of regulation uncertainty that arises from the process of obtaining development entitlements. We combine these sources of market and public sector uncertainty with other key facets of redevelopment projects, such as managerial flexibility, monopolistic competition, and the time to build nature of remediation and construction.

Using data from an ongoing brownfield regeneration project we derive the optimal management policy for the firm, and examine the importance of accurately characterizing the investment environment. We find that uncertainty in the future growth of housing prices has a significant effect on the developer's decision rule, given that remediation and development do not occur instantaneously. Furthermore, our results indicate that the time to build nature of remediation is paramount to understanding the impact of the entitlement regulation process on the investment's value. The need for a lengthly environmental cleanup will limit the effect of regulatory lags on investment value, while uncertainty in the outcome of the regulatory process will significantly effect site value.

Within this brownfield valuation framework we incorporate the fact that development firms have some market power. As a result developer behavior is characterized by an optimal switching model combined with a stochastic control accounting for the sale price for finished units. This hybrid control framework is applicable to many problems in economics, however, it has yet to be fully utilized. This is most likely due to the fact that there does not exist a closed form solution for such models. In Chapter 3 we present a new numerical technique to solve this class of combined optimal switching and stochastic control models. For these problems the value function and optimal control policies will be the solution to a set of Hamilton-Jacobi-Bellman quasi-variational inequalities. Through the use of projection methods to approximate the unknown value function, we transform the problem into an extended vertical non-linear complementarity problem that may be solved using Newton's method. This technique is instrumental in obtaining a solution for the value of brownfield remediation and redevelopment projects in Chapter 2. We further illustrate this method in Chapter 3 using the optimal resource extraction problem of Brennan and

Schwartz [1985]. In the past this model has been solved by eliminating or transforming one of the controls so that the problem contains only a switching or a stochastic control. Using the new numerical technique we show that such simplifications may lead to significant errors in valuing investments.

Chapter 3 also considers the closely related class of models that combine impulse and stochastic controls. The solutions for this class of problems are also described by Hamilton-Jacobi-Bellman quasi-variational inequalities, and as such they too may be represented as extended vertical non-linear complementarity problems through the use of projection methods. This technique for solving combined impulse and stochastic control problems is demonstrated using the portfolio management application considered by Oksendal and Sulem [2002]. We also consider an alternative approach to solving combined impulse and stochastic control problems by redefining the model as one that combines an optimal switching control and a stochastic control. In order to demonstrate this alternative approach we consider the exchange rate control application of Cadenillas and Zapatero [2000].

This dissertation concludes in Chapter 4 with the presentation of a new method for estimating the parameters of continuous time Ito processes using low observation frequency data. A key component in stochastic optimal control modeling is the use to stochastic differential equations to define the dynamics of state variables. However the estimation of parameters for these SDEs is typically troublesome due to the absence of known conditional transition densities for many of the interesting processes. A popular solution to this problem is the quasi-maximum likelihood approach that assumes a normal conditional transition density and approximates the conditional moments by discretizing the process with a first order Euler approximation. However, the Gaussian assumption will only be valid in the case where the observational frequency is relatively high, whereas many economic processes are only observed on a monthly or even a quarterly basis. In these cases Monte Carlo experiments have shown that the normality assumption with moments derived from the Euler approximation will produce inaccurate parameter estimates. A theoretically attractive way to improve this quasi-maximum likelihood estimate is to to improve the approximation of the conditional transition density by making the time step small through the incorporation of additional observations between the known data points. Since these sub-observations are unknown they must be integrated out using the Chapman-Kolmogorov equation.

While this approach is theoretically attractive, a number of significant computational issues have limited its accuracy and efficiency. Most notable is the need to compute a multi-dimensional integral in the Chapman-Kolmogorov equation. Previous work in the area has focused on the use of Monte Carlo integration, while we purpose a more efficient algorithm that utilizes multi-dimensional Gauss Hermite quadrature. The benefit of numerical quadrature over Monte Carlo integration is that the integral may be computed to an arbitrary level of accuracy in a significantly shorter period of time, for low to moderate dimensions. As a result of this improvement in the accuracy and efficiency of the integration, the accuracy of the conditional transition density approximation may be further improved through the application of Richardson extrapolation. Given the accuracy of the numerical integration and the benefits of Richardson extrapolation we are able to achieve the same computational accuracy as the best Monte Carlo based estimators, however we only require a one dimensional integral as opposed to the fifteen dimensional integral required in the Monte Carlo case. The method therefore provides a significant improvement in computational efficiency without any loss of accuracy.

Chapter 2

Valuing Brownfield Remediation and Redevelopment Projects

Brownfields are "abandoned, idled, or underutilized industrial and commercial facilities where expansion or redevelopment is complicated by real or perceived contamination" [Kaiser, 1998]. Within the United States it is estimated that at least 500,000 sites fall into this classification [Meyer and VanLandingham, 2000], though there is evidence to suggest the existence of up to 1,000,000 brownfields [Simons, 1998]. The benefits of successful brownfield regeneration are clear; the reduction of public health threats, revitalization of blighted areas, and decreased development pressure on open space [Singer et al., 2001]. Despite the significant public value of redevelopment, government led remediation efforts remain limited, leaving private developers with a major role in brownfield regeneration. In spite of numerous programs and incentives the participation of firms in brownfield redevelopment projects remains low. Developers have cited the difficulty in selecting profitable regeneration projects as a primary reason for their limited involvement [Meyer and Lyons, 2000].

When analyzing potential investments in brownfields, the current standard is to apply traditional discounted cash flow techniques such as net present value and internal rate of return. However, these methods are incapable of accurately capturing the value of regeneration projects due to the complexity of the remediation and redevelopment process.

The interplay between uncertainty, irreversibility, and managerial flexibility is lost with these simple valuation frameworks. In response, the industry has shown interest in utilizing a real options approach that will facilitate a better understanding about the value of redevelopment projects. In consultation with a leading remediation and redevelopment firm, we design a model that incorporates the characteristics which are seen as critical determinants of brownfield value. These factors include monopolistic competition, public sector and market uncertainty, time to build, and managerial flexibility all within a multi-stage investment framework.

A key feature of real estate projects, often overlooked in the investment literature, is the effect of lags due to regulation processes such as rezoning and permitting. In cases where such lags have been incorporated, the regulation process is unrealistically assumed to be deterministic. In actuality both the length of the regulatory lag and the outcome of the process are associated with uncertainty. Development firms have stated an interest in understanding the impact of this regulation process on the value of sites requiring lengthly environmental remediation prior to development. Using data from an ongoing remediation and redevelopment project, we find that the common practice of ignoring this regulatory process when analyzing brownfield investments will lead to significant valuation errors.

In addition we examine the importance of deviating from the conventional simplifying assumption that housing prices follow geometric Brownian motion. This common conjecture prohibits real estate prices from exhibiting volatility in a form consistent with empirical observations, by virtually eliminating the probability of housing "bubbles". The inclusion of such short term deviations from long term price trends has important implications for the behavior of firms given the temporal nature of both real estate development and environmental remediation. By better describing how developers are making decisions we may increase our understanding of the value of managerial flexibility and in turn provide more appropriate valuations of such investments. Correctly accounting for the time needed to complete lengthy remediation and redevelopment projects, along with incorporating accurate descriptions of real estate price dynamics are paramount in valuing brownfield investments.

The framework developed to capture these features, in addition to updating the literature on brownfield valuation, provides two unique contributions to the real options literature. First, to capture the regulation uncertainty within the multistage investment we introduce the concept of a regulated compound exchange option where obtaining subsequent options requires both endogenous and exogenous actions. Second the existence of monopolistic competition is handled through the application of a hybrid optimal stochastic control model that contains both discrete and continuous controls. This work constitutes a step forward in developing real options based models that realistically depict the investment environment faced by industry. Combining techniques such as time to build, multi-factor stochastic processes, multiple types of regulation uncertainty, in addition to the inclusion of both discrete and continuous control variables, we allow for a more accurate description of such projects.

The remainder of the paper proceeds as follows: Section 2.1 outlines the history of the brownfield problem and briefly discusses the current state of brownfield valuation techniques; Section 2.2 details the model describing the processes of remediation and development; Section 2.3 defines the problem of valuing the investment given the model of redevelopment; Section 2.4 examines implications of the model through an empirical case study; and Section 2.5 contains concluding remarks.

2.1 Brownfield Problem

The current abundance of brownfields is a reaction to both government regulation and changing economic conditions. With enactment of the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) in 1980, any party associated with property containing environmental contamination was subject to possible financial liability, even those that were in no way responsible for the pollution [Bartsch and Collaton, 1995]. At the same time the structural shift in the United States economy away from manufacturing has resulted in downsizing and plant closings leaving abandoned or idle industrial facilities in most communities [Collaton and Bartsch, 1996]. The ambiguous liability and open ended definition of contamination within the legislation resulted in devaluation for hundreds of thousands of such sites, most with little to no contamination [Chilton, 1998].

For developers, a lack of experience in environmental cleanup and the absence of definitive EPA guidelines for evaluating the extent of site contamination limited their ability to accurately estimate the capital expenditures required for remediation [O'Brien, 1989]. In most cases cleanup costs were severely overestimated and as a result brownfield regeneration projects were passed over in favor of greenfield development.[1] In addition to the uncertainty of remediation expenses, the potential liability that the CERCLA attached to third party financial institutions, led most commercial lenders to avoid investments associated with environmental contamination [Meyer and Lyons, 2000]. As a result of the high risk and lack of available capital, redevelopment of brownfields was limited to a small number of sites.

In response to the problems surrounding brownfield redevelopment a number of regulatory changes have been implemented. Most notable is the passage of the Asset Conservation, Lender Liability, and Deposit Insurance Act of 1996, which has provided limited liability relief to financiers of brownfield regeneration projects. To further increase the availability of capital many state and federal programs have begun to offer additional financing for the remediation of environmental contamination prior to development. Furthermore, organizations such as the American Society of Testing and Materials, generated widely accepted guidelines for establishing the extent of contamination on a site [Meyer and VanLandingham, 2000]. The availability of these standards aided in the creation of markets for environmental insurance contracts, and as a result developers became able to reduce their exposure to risk through environmental liability, cost-cap, and prospective liability insurance. As a result of access to capital and the ability to better manage environmental risk private firms with knowledge of remediation techniques found profitable ventures in brownfield regeneration projects. By 2000 private developers had become an integral part of the remediation process, with the three largest firms investing more capital in brownfields than all state governments combined [Meyer and Lyons, 2000].[2] In particular such firms have been responsible for the majority of medium to large scale brownfield cleanup efforts.[3] Despite the success of some developers, many investors choose to either forgo completely or have limited involvement with brownfields due to the difficulty in valuing such projects.

[1] Greenfields are commonly defined as previously undeveloped land that is either open space or farmland.
[2] This excludes state spending on Superfund sites and underground storage tanks.
[3] Medium to large brownfields refers to sites over 5 acres.

Previous attempts to provide a framework for valuing regeneration projects have failed to properly characterize the uncertainty associated with such investments in an approach that may be calibrated using real data.

2.1.1 Previous Approaches to Value Brownfields

The work in Patchin [1988] and Patchin [1991] began the pertinent discussion of how to correctly value property that is affected by either real or perceived environmental contamination. This research examined the importance of including brownfield specific investment characteristics such as remediation costs and indemnity in the valuation of such sites. Incorporating these concepts Mundy [1992b] developed a discrete time discounted cash flow model to determine the value of brownfield redevelopment projects. The approach was to adjust the net present value (NPV) of the final development for cleanup costs associated with the environmental contamination. Additional work by Mundy [1992a] enhanced the approach to provide further realism by accounting for extra investment attributes. Similar research providing further extensions to the discounted cash flow technique of modeling brownfield value may be found in: Chalmers and Roehr [1993], Fisher et al. [1992], Richards [1996], Syms [1996], Tonin [2006], Wilson [1994]. Despite its popularity and the numerous variants, at its core this approach to modeling the value of brownfield regeneration projects overlooks an essential characteristic of such investments by fully ignoring uncertainty.

Prior to the sale of a completed development project, the firm must participate in a lengthy construction process. During this period the market price of the final product, that is real estate prices, will be continuously evolving in a non-deterministic fashion. Paramount to capturing this uncertainty in the investment model is an understanding that as future prices are realized the firm will reevaluate management policies. Therefore, to accurately analyze the value of development projects, careful consideration must be given to the firm's operating flexibility. Traditional discounted cash flow models such as net present value are unable to incorporate this class of optimal control in the face of uncertainty [Dixit and Pindyck, 1994, Trigeorgis, 1993b, Schwartz and Trigeorgis, 2004]. That being the case, it is necessary to utilize a real options approach when modeling the value of real estate development projects [Sirmans, 1997]. What's more, Quigg [1993] and Cunningham [2006] have found empirical evidence to support the use of real options models in pricing real estate investments. The importance of incorporating the value of managerial flexibility is only exacerbated when one considers investments in redevelopment projects, as the firm must first undergo a lengthly remediation phase prior to beginning of construction.

In the literature on applying real options techniques to evaluating real estate investments (Williams [1991]; Capozza and Sick [1991]; Quigg [1995]; Holland et al. [2000]; Capozza and Li [2002]) there has been relatively little research devoted to the problem of valuing brownfield remediation and redevelopment projects. Lentz and Tse [1995] were the first to suggest that discounted cash flow models were unfit to value real estate associated with environmental contamination. In response they developed a model for brownfield regeneration projects where the investment is represented as a compound exchange option. The initial option held is to remediate the site, upon execution the investor receives the option to develop and sell the property.

Paxson [2007] purposed a similar model based on the idea of a compound exchange option. Again the investment is portrayed as a two stage process, where the developer first decontaminates the site and then engages in redevelopment for which the firm receives a lump sum payment upon completion. It is assumed that the expenditure necessary to complete the remediation is associated with some uncertainty. To include this feature the cost of executing the initial option is said to be a process following geometric Brownian motion (GBM). Furthermore, to characterize the uncertainty within the real estate market the payment to the firm for the completed development is also considered to follow GBM. The only significant difference from previous work is the inclusion of uncertainty in the cost of construction, by the assumption that such expenditures are directly proportional to the stochastic cost of remediation. The work of Lentz and Tse [1995] and Paxson [2007] is significant for stating the need to develop a real options approach to brownfield valuation, though the simplified framework ignores important characteristics that are crucial to understanding the value of redevelopment investments. The representation of the project as a set of two simple American options implicitly assumes that remediation and development are instantaneous, thereby ignoring the critical time to build aspect of redevelopment.

Espinoza and Luccioni [2005] draw on the work of Majd and Pindyck [1987] to incorporate this time to build property into the valuation of brownfield redevelopment projects. The process of remediation is described as a continuous operation requiring a constant negative cash flow over a predetermined period of time. Upon completion of the stage the firm receives a one time payment associated with the value of the decontaminated property. This payment is assumed to be uncertain and is described as a GBM process. Once the essential time to build characteristic has been integrated into the model it is imperative that careful consideration be given to the specification of managerial flexibility. To this end, Espinoza and Luccioni [2005] provide the firm with the ability to indefinitely suspend operations during the remediation process. This approach takes an important step in incorporating the time to build characteristic.

We extend this work with three major additions. First, we consider the fact that there still exists uncertainty within the remediation stage despite the predictability of cleanup expenditures and availability of environmental insurance. In order to begin construction, the firm must not only decontaminate the site but also obtain development entitlements from the government. This regulatory process is fraught with uncertainty in both the time it will take to complete and the final outcome. Second, in order to calibrate the uncertainty accompanying the development value it is necessary to model its relationship with prices in the real estate market. Therefore, we carefully model the details of the development process including the underlying dynamics of the market price, the time to build nature of construction, and the firm's ability to influence market absorption. This approach provides a means to calibrate the parameters of the model using empirical data. Third, we note the importance of correctly specifying the sources managerial flexibility, and for that reason incorporate additional options, beyond simple suspension of operations into the model. The specific real options available to a developer are site specific, some examples include the ability to alter the allocation of entitlements between commercial and residential development or to donate the remaining property for tax purposes. To demonstrate the importance of including additional sources flexibility we focus on the commonly held option to sell the remaining undeveloped, and possibly contaminated, property. Given these improvements the framework developed within this paper provides a considerably more realistic model for valuing brownfield remediation and redevelopment projects. In addition it offers an ideal setting in which to consider the impact of the regulation process on investment value and optimal firm behavior.

2.2 Model Description

A typical brownfield regeneration project is comprised of two distinct stages. An initial stage during which the firm remediates the site and obtains the necessary entitlements, and a subsequent stage where the property is developed and sold. This represents an extension to the standard sequential investment models developed by Carr [1988] and Trigeorgis [1993a], which have been applied to wide array of investment projects.[4] Traditionally multistage investment models have focused on the case whereby foregoing a specific cash outlay the firm is able to advance to a subsequent stage of the investment. The case of brownfield remediation and redevelopment differs in that to gain the option to develop the site, not only must the firm remove the existing environmental contamination but the regulatory uncertainty must also be resolved. To the best of our knowledge this description of a multistage investment as a compound option in which the investor's ability to obtain the second set of options requires both endogenous and exogenous actions to occur is a unique contribution of this paper. This section proceeds by laying out the model describing the investment environment beginning with the initial remediation stage.

2.2.1 Remediation and Entitlement Stage

During the investment's initial stage the developer is required to both remediate the site and obtain entitlements in order to facilitate future development. Therefore this first stage is denoted as the remediation and entitlement stage. As noted by Mayer and Somerville [2000a] regulation lags mainly occur at two major stages in the development process. An initial delay in order to obtain zoning and subdivision approval and a subsequent process of obtaining building permits. The permitting procedure tends to be relatively quick and a successful outcome is typically certain. Therefore we choose to ignore this lag

[4] For a brief introduction to the types of applications using option pricing theory to evaluate multistage investment projects we refer the reader to Panayi and Trigeorgis [1998], Alvarez and Stenbacka [2001], Lee and Paxson [2001], and Rogers et al. [2002].

in favor of focusing on the more lengthy initial regulatory process in which local authorities will often require changes in the project's design such as development density. The stage is considered to begin with a request for entitlements to develop ϵ square feet of residential housing. After the preliminary application is submitted to the appropriate government authority, the process continues to require action on behalf of the firm. Therefore there exists a negative cash flow C_E for various soft costs, such as legal fees, until the completion of the regulatory process. To the developer the time until final approval is both exogenous and uncertain. This characteristic is captured by describing the time until conclusion of the regulatory review as a random variable from the exponential distribution, with mean λ.[5]

To a developer the final outcome of the regulatory process is also considered to be uncertain and exogenous to their actions. When project approval is granted by the local authorities it is modeled as the proportion of the initial application amount ϵ for which the firm will be entitled. This proportion will be defined as $\pi \geq 0$, such that the firm will be entitled to develop $\pi\epsilon$ square feet of residential housing. Due to the uncertainty in the regulatory proceedings π is defined as a random draw from a discrete distribution of N possible outcomes where $\pi_1 = 0$ and $\sum_{j=1}^{N} Pr(\pi_j) = 1$. We introduce the variable Y_t to describe the state of the firm's development entitlements, such that

$$Y_t = \begin{cases} 0 & \text{if approval has not been received} \\ j & \text{if approval has been received for } \pi_j\epsilon \text{ units} \end{cases},$$

where $Y_t \in \{0, 1, \ldots, N\} =: \mathcal{Y}$.

In addition to acquiring development entitlements the firm must also remediate the site during the initial stage of the investment. These actions may take place simultaneously or the developer may choose to delay remediation while observing the entitlement process along with changes in the housing market. Based on initial evaluations of site contamination the firm is assumed to know the total amount of remediation required. This is defined by the total time, $\bar{R}$, the firm must be actively cleaning the property in order to complete the process. In other words at the beginning of the project, $t = 0$, the remediation time

[5]This assumption is analogous to describing entitlement regulation as a Poisson process with mean arrival rate $1/\lambda$.

remaining is $R_0 = \bar{R}$. The dynamics of this state variable are dependent upon whether the firm is actively remediating,

$$dR_t = \begin{cases} -dt & \text{if actively remediating} \\ 0 & \text{otherwise} \end{cases}.$$

While actively undertaking remediation the firm incurs a negative cash flow C_H associated with the hard costs of removing the environmental contamination and C_R associated with soft costs such as management and marketing fees.[6] This deterministic description of the remediation process represents the ability of experienced brownfield developers to use their knowledge of contamination assessment and state of the art remediation technologies, to accurately evaluate the cost and time required for the process. Well functioning markets for environmental insurance allow for firms to hedge against unforeseen contamination further strengthening the assumption of a deterministic cost for remediation of the site.[7]

If the firm is not currently remediating, they still face a negative cash flow C_N associated with soft costs such as management fees and maintaining security on the site. The overall cash flow for the firm during the remediation and entitlement stage, will depend on the current state of the regulatory process, Y_t, along with the actions of the firm. That is if the firm is actively remediating and has not yet received entitlement approval ($Y_t = 0$) the cash flow will be $-(C_H + C_R + C_E)$ as opposed to $-(C_N + C_E)$ if the firm has suspended remediation. If the regulatory process has already been completed ($Y_t > 0$) the cash flow while the firm is active or suspended will reduce to $-(C_H + C_R)$ and $-C_N$ respectively.

The cumulative cash flow during the initial stage of the investment will always be negative and may be viewed as the expenditure necessary to acquire the option to develop

[6]We make this distinction due to the fact that certain expenditures during this stage of the project are unique to the firm and its intended use of the remediated property. Therefore not all expenditures will impact the resale value of the contaminated site if the firm chooses to forego the remainder of the project. Those costs that may viewed as a financial liability acquired by the firm at the time of the site's purchase are contained within C_H.

[7]It may be the case that the developer anticipates a significant probability that cost overruns will surpass the cost-cap insurance for the project. In response the hard costs associated with remediation, C_H, may be described as a continuous time random variable that may increase with some positive probability. Alternatively, it may be the case that the remediation time remaining, R_t, varies randomly. In consultation with our corporate partner we have chosen to rule out this possibility by defining C_H as a constant and R_t as deterministic, but note that such environmental risk may be incorporated in the modeling framework presented.

the property. Though it may be the case that during the life of the investment the expected remaining expenditures may no longer be warranted given the state of the market and the expected value of the developed property. In these circumstances it may be optimal for the developer to pursue other uses of the property. This managerial flexibility is analogous to the firm holding a set of real options. In this paper we focus specifically on the firm's option to sell the site "as is" to another party, an option common to most brownfield redevelopment projects. We note that the framework being developed may easily be augmented to include additional site specific options, for example the ability to donate the site as a park for tax purposes or rezoning the property for commercial use.

In order to describe the payoff from exercising the option to sell the site in its current state, we assume that the value of the land is derived from the value of its possible development use, adjusted for the remaining remediation costs. Additionally it is assumed that the value of undeveloped property is directly proportional to the market value of developed real estate, P_t. Therefore the salvage value for the firm is represented by

$$\eta P_t U(Y_t)\epsilon - C_H R_t, \tag{2.1}$$

where

$$U(Y_t) = \begin{cases} E[\pi] & \text{if } Y_t = 0 \\ \pi_{Y_t} & \text{if } Y_t > 0 \end{cases}, \tag{2.2}$$

and η is the proportion of the price for developed real estate attributed to the value of the land. As a result, the first part of (2.1) is interpreted as the current value of the land that may eventually be developed. In the case where entitlements have not yet been received this value is an expectation based on the distribution governing the possible outcomes of the approval process. The second component of the salvage value is the cost of the remediation that is still required before development may begin.

2.2.2 Development Stage

Upon completion of the remediation and entitlement stage, that is once the firm has completed the regulatory process and removed the environmental contamination, they hold the option to begin development of the site. The value of the development stage will be dependent upon the state of the market for residential housing. It is assumed that the market price for developed real estate, P_t, is governed by an Ito process of the form

$$dS_t = \mu(S_t)dt + \sigma(S_t)dW_t, \tag{2.3}$$

where $S_t \in \mathbb{R}^d$ is a set of d state variables including the market price P_t; $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the drift function and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ is the diffusion function. W_t is a d-dimensional Brownian motion in the $(\Omega, \mathcal{F}, \mathbb{P})$ probability space with the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. To ensure the existence of a unique solution to (2.3) it is assumed that $\mu(\cdot)$ and $\sigma(\cdot)$ are Lipschitz continuous and that

$$|\mu(S_t)|^2 + ||\sigma(S_t)||^2 \leq Q(1 + |S_t|^2)$$

for some constant Q, where $|\cdot|$ and $||\cdot||$ are the vector and matrix norm respectively, therefore ruling out explosive growth [Oksendal and Karsten, 1998]. At this stage no further assumptions are made about the structure of the drift and diffusion in order to emphasize the generality of the framework. We note that the particular form of the process appropriate for a project will be site specific due to the spatial uniqueness of real estate markets.

We assume that the developer is operating under monopolistic competition and therefore holds the option to adjust its output price.[8] In particular the price at which the developer sells completed units is modeled as proportional to the market price. That is the sale price will be θP_t, where $\theta \in \mathbb{R}_+$ represents a continuous control for the firm.

The rate at which the market will absorb the firm's development, A, will be dependent upon the sale price so that $A = A(\theta)$. In essence the function defining the absorption rate is the demand curve for the firm's completed housing. As such it will be downward sloping in the control θ, which represents the relative price for the firm's product. For simplicity it is assumed that the absorption rate will be a linear function of the relative price,

[8] The assumption of monopolistic competition was included after discussions with the remediation firm consulted for this project. For many urban areas the majority of new development available to home buyers is on the fringe, limiting those interested in closer proximity to the city center to preexisting housing [Mayer and Somerville, 2000b]. Therefore developers of the large to medium brownfields commonly found within urban areas have the ability to provide additional amenities beyond those common to the preexisting real estate and as a result the firm experiences some market power.

θ, on the range $[\Theta_L, \Theta_H]$, where Θ_H represents the price markup beyond which consumer demand is zero, and Θ_L corresponds to an absorption rate equal to the firm's maximum construction rate, $\bar{\kappa}$. Therefore the function defining the absorption rate for the firm's product, based on their choice of sale price, is piecewise linear of the form

$$A(\theta) = \begin{cases} \bar{\kappa} & \text{if } \theta \leq \Theta_L \\ \kappa_0 + \kappa_1 \theta & \text{if } \Theta_L < \theta < \Theta_H \\ 0 & \text{if } \theta \geq \Theta_H \end{cases},$$

where κ_0 and κ_1 are the intercept and slope of the linear function respectively. These parameters are defined as

$$\kappa_0 = \frac{\bar{\kappa}\Theta_H}{\Theta_H - \Theta_L}$$

and

$$\kappa_1 = \frac{-\bar{\kappa}}{\Theta_H - \Theta_L}.$$

Figure 2.1 illustrates this function.

It is assumed that the developer will operate in a manner consistent with maintaining a zero inventory, that is the rate of construction will be equal to rate of absorption. This represents a continuous analog to the case in which the firm develops the site in phases and sells each one upon completion. Therefore the dynamics for the remaining number of units the firm is entitled to construct, K_t, may be defined as

$$dK_t = -A(\theta)dt,$$

where the initial level was determined by the outcome of the regulatory process, $K_0 = \pi_Y \epsilon$.[9] The marginal cost of development, C_A, is considered to be constant over time.

While the firm is actively developing the site they are subject to a flow of expenditures C_M associated with soft costs such as management fees, marketing, and onsite overhead. Therefore the net cash flow for to the firm when engaged in active construction

[9]The zero inventory assumption represents a simplification that was made in consultation with our corporate partner, who considered this to be a minor issue. The benefit is that the dimensionality of the problem may be reduced by eliminating the need to keep track of the current inventory of unsold units.

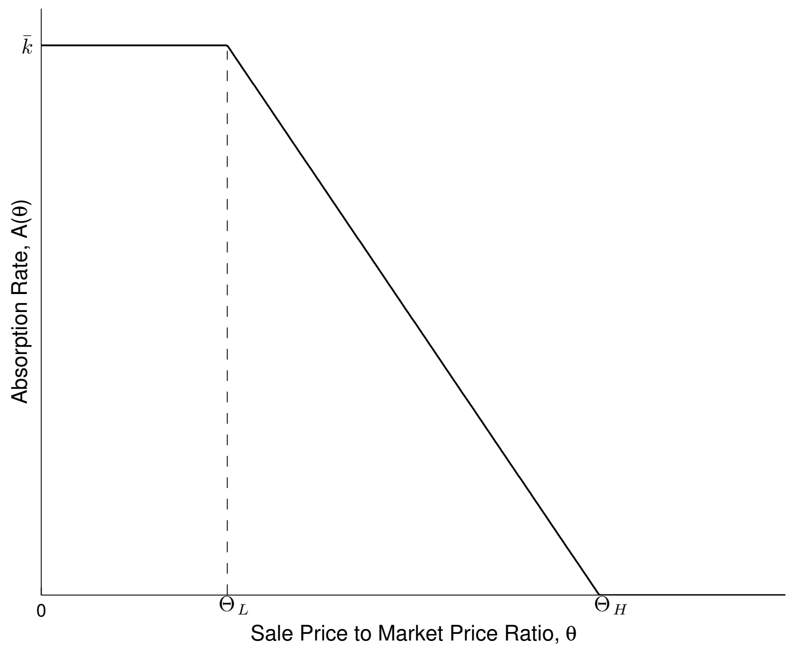

Figure 2.1: Piecewise Linear Absorption Rate

has the form

$$A(\theta)\,(\theta P_t - C_A) - C_M.$$

We assume that the firm has the option to suspend construction of new units, $A = 0$, and as a result faces a lower cost C_S, such that $C_S < C_M$. Therefore the net cash flow for the firm when operations are suspended is $-C_S$.

In addition to the firm's ability to optimally control the sale price and suspend operations, it is important to include other sources of managerial flexibility in order to fully capture the development value of the site. As previously discussed, the specific real options held by the developer will be dependent upon both the site and the firm's contractual obligations. In this model we consider the option to sell the remaining property, and continue with the same assumptions as in the remediation and entitlement stage. That is the salvage value is determined by the current market value of the entitled construction. Therefore by exercising the option to sell the remaining undeveloped property the firm receives a one time payment of $\eta P_t K_t$ in exchange for the right to future development.

2.3 Investment Valuation

The objective of the firm is to maximize, with respect to its operating choices, profits over an infinite time horizon.[10] Therefore determining the investment's value will require an understanding of the optimal control solution to the dynamic optimization problem associated with the firm's actions. This section outlines the conditions that define both the value of the project and the developer's optimal behavior. Despite the dependence of the development stage on the outcome of the entitlement process, its value may be determined independently over the whole range of possible values for the remaining construction, $[0, \pi_N \epsilon]$. This solution may then be used to assign the terminal value of the remediation stage. Given its importance in valuing the project during the initial remediation and entitlement stage we begin by discussing the problem of valuing the development stage. Afterwards we present the problem of valuing the remediation and entitlement stage given the development value.

2.3.1 Value of the Development Stage

The development stage may be characterized as a combined optimal switching and continuous control model, where the investment is considered to be in one of three regimes corresponding to active construction, suspension of operations, and sale of remaining property. The current regime in time t is denoted by Z_t such that $Z_t \in \{1, 2, 3\} =: \mathcal{Z}$ where

$$Z_t = \begin{cases} 1 & \text{if the remaining property has been sold} \\ 2 & \text{if construction is currently suspended} \\ 3 & \text{if the firm is actively developing the site} \end{cases} .$$

The firm's choice of the current regime will be denoted as a discrete control as it may only take on one of countably many values.

In addition to the choice of regime, the developer has control over the sale to market price ratio θ_t when construction is active, $Z_t = 3$. The variable θ_t will be denoted as a continuous control as opposed to the discrete control Z_t. It is assumed that with respect

[10] In practice there may be time constraints on the project, but including such characteristics would increase the dimensionality of the problem and in turn make solving the model considerably more difficult.

to the filtered probability space of the stochastic process, S_t as defined by (2.3), that

$$\theta_t \text{ is } \mathcal{F}_t\text{-measurable for all } t \geq 0. \tag{2.4}$$

A policy w for the discrete control variable may be described as the possibly finite double sequence

$$w = (\tau_1, \tau_2, \ldots, \tau_i, \ldots; \zeta_1, \zeta_2, \ldots, \zeta_i, \ldots),$$

where $0 \leq \tau_1 \leq \tau_2 \leq \ldots$ are stopping times with respect to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. Associated with the stopping times are the changes to the discrete control $\zeta_i \in \mathcal{Z}$. In other words at time τ_i the agent switches to the regime ζ_i. Therefore a policy for the combined continuous and discrete control may be written as $\nu = (\theta, w)$. The combined policy ν is considered to be admissible if (2.4) holds, $\theta \in [\Theta_L, \Theta_H]$, $\zeta_i \in \mathcal{Z}$ $\forall i$, and τ_i are stopping times. The set of all admissible combined controls is denoted as V.

When the system is in the particular state (s, k, z), with the continuous control θ the firm receives the net cash flow $g(s, k, z, \theta)$ where

$$g(s, k, z, \theta) = \begin{cases} 0 & \text{if } z = 1 \\ -C_S & \text{if } z = 2 \\ A(\theta)\left(\theta P_t - C_A\right) - C_M & \text{if } z = 3 \end{cases},$$

recalling that $P_t \in S_t$. We refer the reader to Table 2.1 for a summary of the notation. Given the state (s, k, z) the cost of switching from the current regime z to ζ is denoted by $H^D(s, k, z, \zeta)$ where

$$H^D(s, k, z, \zeta) = \begin{cases} -\eta p k & \text{if } z > 1 \text{ and } \zeta = 1 \\ \infty & \text{if } z = 1 \text{ and } \zeta \neq 1 \\ 0 & \text{otherwise} \end{cases}.$$

This definition states that the act of suspending and (re)starting active construction is costless, and that by selling the remaining property the developer receives a one time payment

ηpk.[11] The implementation of an infinite switching cost is analogous to the assumption that once the firm sells the property they are unable to regain ownership in the future. The expected discounted flow of benefits to the agent at the onset of the stage and under the policy $\nu \in V$ may then be defined as

$$\begin{aligned} G^{\nu}(s,k,z) = \quad & E\left[\int_0^\infty e^{-\rho t} g(S_t, K_t, Z_t t, \theta_t)dt \right. \\ & \left. - \sum_{j=1}^\infty e^{-\rho\tau_j} H^D(S_{\tau_j}, K_{\tau_j}, \zeta_{j-1}, \zeta_j) \right| S_0 = s, K_0 = k, Z_0 = z \Bigg], \end{aligned}$$

where ρ is the risk-adjusted discount rate and $Z_t = \zeta_j$ if $\tau_j \leq t < \tau_{j+1}$. It may be seen that if $z = 1$, then $G^{\nu}(s,k,z) = 0$ due to the constraint that the firm may not reacquire the project once it has been sold. We therefore restrict the discussion to the meaningful problem of finding, for all (s,k,z) where $z > 1$, the value function for the development stage $V^D(s,k,z)$ such that

$$V^D(s,k,z) = \sup_{\nu \in V} G^{\nu}(s,k,z). \tag{2.5}$$

Due to the absence of a fixed cost associated with switching between the active and suspension regimes it will be the case that $V^D(s,k,2) = V^D(s,k,3)$. Given that the project's value, prior to being sold off, will be independent of the regime we redefine the function as $V^D(s,k) = V^D(s,k,2) = V^D(s,k,3)$.

As a natural extension to the work of Brekke and Oksendal [1994] a viscosity solution for the value function will be defined by the system

$$\rho V^D(s,k) \geq A(\theta^*)\left(\theta^* P_t - C_A\right) - C_M - A(\theta^*)\frac{\partial V^D}{\partial k} + \mathcal{L}V^D(s,k), \tag{2.6}$$

$$\rho V^D(s,k) \geq -C_S + \mathcal{L}V^D(s,k), \tag{2.7}$$

and

$$V^D(s,k) \geq \eta pk, \tag{2.8}$$

[11] It is possible to generalize the framework to include fixed switching costs for suspending and (re)starting active construction. For the case study presented in this paper the developer considered such operating changes to lack any fixed costs. Therefore, the generalization is forgone in favor of a notation simplification.

Table 2.1: Brownfield Redevelopment Model Parameters and Variables

All monetary values are in 2007 U.S. dollars.

Parameter	**Description**	**Value**
L	Cost of brownfield site	47,300,000
ϵ	Requested entitlements (ft^2)	2,601,000
λ	Expected length of entitlement process (years)	0.5
C_E	Annual negative cash flow for the entitlement process	500,000
$\bar{R}$	Total time needed for remediation (years)	2
C_N	Annual negative cash flow while remediation is suspended	886,225
C_H	Annual negative cash flow associated with remediation	2,956,500
C_R	Annual negative soft cost cash flow while actively remediating	2,532,074
η	Proportion of developed value attribute to the land	0.16
Θ_L	Sale to market price ratio below which demand can't be met	0.75
Θ_H	Sale to market price ratio above which demand is zero	1.25
$\bar{\kappa}$	Maximum construction rate (ft^2/year)	1,500,000
C_A	Marginal cost of construction (\$/ft^2)	92
C_M	Annual negative soft cost cash flow while actively developing	2,532,074
C_S	Annual negative cash flow while development is suspended	886,225

Variable	**Description**
S	Vector of d state variable describing the real estate market
P	Market price of completed units (\$/ft^2, component of S)
μ	Growth rate variable (component of S)
R	Remediation remaining (years)
K	Construction remaining (ft^2)
π	Proportion of the entitlement application approved
Y	State of the entitlement process
θ	Sale to market price ratio
A	Absorption rate (ft^2/year)
V^R	Value of the remediation stage
V^D	Value of the development stage

where θ^* is the optimal relative price and $\mathcal{L}$ is the differential generator

$$\mathcal{L} = \sum_{i=1}^{d} \mu_i(s) \frac{\partial}{\partial s_i} + \frac{1}{2} \sum_{i=1}^{d} \sum_{j=1}^{d} \left[\sigma(s)\sigma(s)^T\right]_{ij} \frac{\partial^2}{\partial s_i \partial s_j}. \tag{2.9}$$

In addition one of the conditions in (2.6)-(2.8) must hold with equality. For each (s,k) the equality will determine the firm's optimal management policy with respect to the discrete control. If (2.6) is the condition to hold with equality then it will be optimal at (s,k) for the agent to be actively developing the site. On the other hand if (2.7) holds with equality, then it is optimal for the firm to suspend operations. Alternatively, if it is optimal for the developers to sell off the remainder of the site (2.8) will hold with equality. The value function must also satisfy the boundary condition

$$V^D(s,0) = 0, \tag{2.10}$$

which states that when the firm completes development of the site, $k = 0$, the project no longer holds value for the developer.

The optimal policy for the continuous control, that is the relative price, will be determined by maximizing the value function with respect to $\theta \in [\Theta_L, \Theta_H)$ given that the firm is actively developing, $z = 3$. The fact that it is never optimal to choose a relative price of Θ_H is evident from the conditions in (2.6)-(2.8). If the firm is actively developing the site, it must be the case that (2.6) holds with equality. Combining this result with the condition in (2.7), it will be the case that

$$A(\theta^*)\left(\theta^* p - C_A\right) - C_M - A(\theta^*) \frac{\partial V^D}{\partial k} \geq -C_S.$$

If it were optimal for the firm to choose a relative price at the upper kink in the absorption rate function, $\theta^* = \Theta_H$, so that $A(\theta^*) = 0$, the above condition would reduce to $C_M \leq C_S$. Whereas it was explicitly assumed in Section 2.2 that the fixed cost flow when suspended was less than that of active development, $C_M > C_S$. Therefore, θ^* will exist on the range $[\Theta_L, \Theta_H)$.

Again we note that if it is optimal for the firm to be in the regime associated with active development, $z = 3$, it must the case that (2.6) holds with strict equality. Therefore maximizing the value function with respect to θ may be seen as analogous to the problem

$$\sup_{\theta \in [\Theta_L, \Theta_H)} \left\{ A(\theta) \left(\theta P_t - C_A\right) - C_M - A(\theta) \frac{\partial V^D}{\partial k} + \mathcal{L} V^D(s,k) \right\}.$$

Given that the upper bound on θ is non-binding the Karush-Kuhn-Tucker (KKT) conditions associated with this constrained maximization problem are

$$\frac{\partial A}{\partial \theta} \left(\theta p - C_A\right) + A(\theta) p - \frac{\partial A}{\partial \theta} \frac{\partial V^D}{\partial k} + \upsilon = 0,$$

$(\theta - \Theta_L)\upsilon = 0$, and $\upsilon \geq 0$.

On the interior where $\upsilon = 0$ it will be the case that $A(\theta) = \kappa_0 + \kappa_1 \theta$ and

$$\frac{\partial A}{\partial \theta} = \kappa_1,$$

and therefore the KKT conditions imply that the optimal policy for the continuous control in the relative price will satisfy

$$\kappa_1 \left(\theta^* p - C_A\right) + (\kappa_0 + \kappa_1 \theta^*) p - \kappa_1 \frac{\partial V^D}{\partial k} =,$$

such that the optimal policy θ^* may be written as

$$\theta^* = \frac{1}{2} \left(\frac{C_A + \frac{\partial V}{\partial k}}{p} - \frac{\kappa_0}{\kappa_1} \right). \tag{2.11}$$

We note that this dependent upon an interior solution and that the optimal relative price will not be lower than Θ_L.

A closed form solution for the value function $V^D(s,k)$ defined by the variational inequality in (2.6)-(2.8) and the conditions (2.10) and (2.11) does not exist. We therefore use the numerical developed in Chapter 3 to obtain an approximation of the value function using projection methods and collocation. To handle the dependence of the optimal continuous control policy on the partial derivative of the value function with respect to the remaining construction, θ^* is computed simultaneously with $V^D(s,k)$. It is worth noting that we reduce the dimensionality of the problem by approximating the derivative of the value function with respect to the deterministic state variable K_t using a explicit forward finite difference, and then solve for value function at each point on a discrete grid for K_t. Full details of the numerical method are located in Appendix A.

2.3.2 Value of the Remediation and Entitlement Stage

Similar to the development stage a brownfield regeneration project in the remediation and entitlement stage may be described as an optimal switching model. In the remediation stage the firm is not concerned with setting a sale price and therefore is not considered to have a continuous control. Instead the problem is a pure optimal switching model with a discrete control X_t representing the current regime in time t. Similar to the development stage it will be the case that $X_t \in \{1,2,3\} =: \mathcal{X}$ where

$$X_t = \begin{cases} 1 & \text{if the property has been sold} \\ 2 & \text{if remediation is currently suspended} \\ 3 & \text{if the firm is actively remediating the site} \end{cases} .$$

A policy w for this discrete switching control may again be described as the possibly finite double sequence

$$w = (\tau_1, \tau_2, \ldots, \tau_i, \ldots; \xi_1, \xi_2, \ldots, \xi_i, \ldots),$$

where $0 \leq \tau_1 \leq \tau_2 \leq \ldots$ are stopping times associated with changes to the discrete control $\xi_i \in \mathcal{X}$. The set of all admissible controls is denoted by $\mathcal{W}$.

The state of the system is defined by the variables (S_t, R_t, Y_t, X_t), where S_t represents the state of the real estate market, R_t is the remediation remaining, Y_t is the state of the entitlement process, and X_t is the current regime. When the system is in the state

(s,r,y,x) the net cash flow to the firm is $f(y,x)$ where

$$f(y,x) = \begin{cases} 0 & \text{if project has been sold, } x = 1 \\ -(C_N + C_E) & \text{if suspended and waiting for approval, } x = 2, y = 0 \\ -C_N & \text{if suspended and entitlements approved, } x = 2, y > 0 \\ -(C_H + C_R + C_E) & \text{if active and waiting for approval, } x = 3, y = 0 \\ -(C_H + C_R) & \text{if active and entitlements approved, } x = 3, y > 0 \end{cases} .$$

Given the state (s,r,y,x) the cost of switching from the current regime x to ξ is denoted by $H^R(s,r,y,x,\xi)$ where

$$H^R(s,r,y,x,\xi) = \begin{cases} -\eta p U(y)\epsilon + C_H r & \text{if } x > 1 \text{ and } \xi = 1 \\ \infty & \text{if } x = 1 \text{ and } \xi \neq 1 \\ 0 & \text{otherwise} \end{cases} ,$$

and where U is defined in (2.2). This states that there are no fixed costs associated with (re)starting and suspending remediation and that the lump sum payment received from selling the property is $\eta p U(y)\epsilon - C_H r$, consistent with the definition in (2.1). As with the development stage, it is assumed that once the developer sells the property there exists an infinite cost to reentering the project. Given the firm's cash flow and switching costs the expected discounted flow of profits to the developers at the onset of the stage and under the policy $w \in \mathcal{W}$ is defined as

$$\begin{aligned} F^w(s,r,y,x) = \quad & E[\textstyle\int_0^\infty e^{-\rho t} f(Y_t, X_t)dt \\ & - \textstyle\sum_{j=1}^\infty e^{-\rho\tau_j} H^R(S_t, R_t, Y_t, \xi_{j-1}, \xi_j) \Big| S_0 = s, R_0 = r, Y_0 = y, X_0 = x], \end{aligned}$$

where ρ is the risk adjusted discount rate and $X_t = \xi_j$ if $\tau_j \leq t < \tau_j + 1$. As was the case in the development stage, it may be seen that when the firm has sold the project, $x = 1$, the expected discounted profit flow is equal to zero, $F^w(s,r,y,1) = 0$. The problem is then to find for all (s,r,y,x) where $x > 1$ the value function for the development stage $V^R(s,r,y,x)$ such that

$$V^D(s,r,y,x) = \sup_{w \in \mathcal{W}} F^w(s,r,y,x). \tag{2.12}$$

The absence of a fixed cost associated with switching between the suspension and active remediation regimes implies that $V^R(s,r,y,2) = V^R(s,r,y,3)$. Therefore in order to simplify the notation we redefine the value function as $V^R(s,r,y) = V^R(s,r,y,2) = V^R(s,r,y,3)$.

Similar to the development value we define the value of the remediation and entitlement stage as the solution to a set of variational inequalities. The system of conditions differs from those of Section 2.3.1 due to the lack of a continuous control and the inclusion of the regulatory process. The value function $V^R(s,r,y)$ will therefore satisfy the conditions

$$\rho V^R(s,r,y) \geq f(y,3) - \frac{\partial V^R}{\partial r} + \mathcal{L}V^R(s,r,y) + \frac{1}{\lambda}\left\{E\left[V^R(s,r,\tilde{y})\right] - V^R(s,r,y)\right\}, \quad (2.13)$$

$$\rho V^R(s,r,y) \geq f(y,2) + \mathcal{L}V^R(s,r,y) + \frac{1}{\lambda}\left\{E\left[V^R(s,r,\tilde{y})\right] - V^R(s,r,y)\right\}, \quad (2.14)$$

and

$$V^R(s,r,y) \geq \eta p U(y)\epsilon - C_H r, \quad (2.15)$$

where $\mathcal{L}$ is the differential generator as defined in (2.9) and $\tilde{y}$ represents the value of Y_t after approval has been received. Since this value is a random variable the expected change in the value function from the entitlement process, $\frac{1}{\lambda}\left\{E\left[V^R(s,r,\tilde{y})\right] - V^R(s,r,y)\right\}$, is defined in terms of the expectation of the value function after approval. The distribution governing $\tilde{y}$ is assumed to be discrete and therefore the expectation of the value function may be defined as

$$E\left[V^R(s,r,\tilde{y})\right] = \sum_{j=1}^{N} Pr(\pi_j)V^R(s,r,j). \quad (2.16)$$

In the case where regulatory process has already concluded $y > 0$, we interpret $\tilde{y}$ as y so that
$\frac{1}{\lambda}\left\{E\left[V^R(s,r,\tilde{y})\right] - V^R(s,r,y)\right\} = 0.$

One of the conditions in (2.13)-(2.15) must hold with equality. For each (s,r,y) the optimal control policy for the firm will be determined by which condition holds as a strict equality. If it is the case that (2.13) holds with strict equality then it is optimal at (s,r,y) for the firm to actively remediate the site. However, if (2.14) were to hold with strict equality instead, then it is optimal for the firm to suspend remediation. If it is (2.15) that holds with equality then the firm will sell off the project.

When the firm has completed remediation, $r = 0$, and if the regulation process has been concluded, $y > 0$, then the developer will move on to the second stage of the investment. Therefore the value function will be equal to that of the development stage, such that

$$V^R(s, 0, y) = V^D(s, \pi_y \epsilon), \qquad y > 0. \tag{2.17}$$

In the case where the firm has completed remediation, $r = 0$, but has yet to receive development entitlements, $y = 0$, the firm still has the ability to sell the project and therefore the value will be the solution to an optimal stopping problem with an exogenous shock (i.e. regulatory approval) that moves the project into the development stage. This problem is similar to that of the remediation stage prior to regulatory approval except that the firm only holds one option, and that is the ability to sell the property. As a result the value $V^R(s, 0, 0)$ will be the solution to the variational inequality defined by

$$\rho V^R(s, 0, 0) \geq -(C_E + C_N) + \mathcal{L}V^R(s, 0, 0) + \frac{1}{\lambda}\left\{E\left[V^D(s, \pi\epsilon)\right] - V^R(s, 0, 0)\right\} \tag{2.18}$$

and

$$V^R(s, 0, 0) \geq \eta p E[\pi]\epsilon, \tag{2.19}$$

where $\mathcal{L}$ is the differential generator in (2.9) and the expectation of the development value with respect to the entitlement outcome is defined as

$$E\left[V^D(s, \pi\epsilon)\right] = \sum_{j=1}^{N} Pr(\pi_j) V^D(s, \pi_j \epsilon).$$

One of the conditions in (2.18)-(2.19) must hold with equality. If (2.19) holds with equality than it is optimal for the developer to sell the project, otherwise the firm will continue to wait for the regulatory decision.

The value function for the remediation and entitlement stage $V^R(s, r, y)$ is defined by the variational inequality in (2.13)-(2.15) along with the boundary conditions in (2.17)-(2.19). A closed form solution for the value function does not exist though a numerical approximation may be obtained through the use of finite difference and projection techniques. In order obtain a more tractable problem we appeal to the fact that once the regulation process has come to a conclusion the development entitlements are fixed for the remainder of the project. Therefore the value of the remediation stage after the regulation process has ended may be determined independently for each possible outcome. These values may then be employed in the computation of (2.13). For each of these $N+1$ switching models the dimensionality of the problem is further reduced by approximating the derivative of the value function with respect to the deterministic state variable R_t using a explicit forward finite difference. At each time step the value functions are then approximated using a projection method with collocation. Full details of the numerical method are included in Appendix B.

2.4 Empirical Case Study

In order to test the valuation model presented and assess the impact of the regulation procedure we consider an empirical case study using a site recently acquired by our corporate partner. The property is a 402 acre parcel within the Houston / Sugar Land / Baytown metropolitan statistical area (MSA) with a cost to the firm of $47,300,000. This site was previously used for an industrial application which resulted in soil and groundwater contamination. The primary source of contamination is the existence of petroleum hydrocarbons in the soil, due to leaking underground storage tanks for fuel and oil. The extended presence of the contaminants has resulted in groundwater pollution which requires extensive remediation and monitoring. In addition lead and arsenic exist in the soil as a result of a firing range and railroad previously located on the site. Also considered part of the remediation process is the demolition of any structures that will not be redeveloped, thus requiring the removal of an estimated 27,800 linear feet and 54,700 square feet of asbestos. After remediation of the property the firm intends to construct primarily residential housing on the site. This section proceeds by presenting the specifics of the investment, followed by a discussion on the dynamics of real estate prices. We conclude the section with results from the empirical case study.

2.4.1 Parameter Calibration

Based on the extent of the environmental contamination in the soil and the volume of asbestos it is estimated that the total cost of remediation to the firm will be \$5,613,000. The expectation is that the cleanup and demolition of the site will take two years. In order to limit exposure to the risk associated with the contamination a \$10 million insurance policy for environmental and prospective liability will be purchased at a cost of \$300,000. Therefore the average annual negative cash flow associated with the remediation process is estimated to be \$2,956,500.

The firm expects that soft costs such as corporate management, marketing and on-site overhead will total \$15,192,446 over the life of the investment, which is assumed to be six years. This results in an average soft cost flow of \$2,532,074 per year while the investment is in an active state. If the firm has suspended operations it is assumed that it will only have to cover corporate management fees along with typical mothballing costs such as on-site security. This is assumed to be 35% of the active soft costs or \$886,225 per year. There is assumed to be no difference in the soft costs between the remediation and development stage.

During the initial stage of the investment the our corporate partner plans to request entitlements to construct 867 residential units with an average floor plan of 3,000 square feet for a total of 2,601,000 square feet of development. Conditional on prior information regarding the local government the firm expects the entitlement process to last 6 months. This corresponds to mean arrival time of 1/2 for the random variable governing the timing of the regulation phase. The discrete distribution for outcome the entitlement process has been parameterized by our corporate partner based on their previous experience and initial discussions with the municipality. The 8 possible outcomes and the corresponding probabilities are presented in Table 2.2. It is worth noting that there is a possibility for the firm to be granted entitlements beyond what was applied for, $\pi = 1.1$. While engaged in the regulation process the firm will face an average annual negative cash flow of \$500,000 for management and legal fees.

Once in the development stage the firm faces an average construction cost of \$92

Table 2.2: Discrete Distribution for Entitlement Process Outcome

π	$Pr[\pi]$
0	.02
.5	.04
.6	.07
.7	.10
.8	.20
.9	.40
1	.15
1.1	.02

per square of housing. This value includes the hard costs associated with the construction of the residential housing along with non-reimbursable infrastructure expenses. The per square foot sale price of development is proportional to the market price as defined by the choice of θ. This continuous control dictates the development stage cash flow through the per unit net income and the absorption rate. The piecewise linear function describing the relationship between the relative price and the absorption rate, represents the most difficult component to calibrate. After consulting with out corporate partner we set the maximum construction rate to be 1,500,000 square feet per year. At a relative price of .75 it is assumed that the market absorption would reach this maximum construction rate, and at a markup of 1.25 demand for the firm's housing would be equal to zero.

As previously discussed the firm holds the option to sell the undeveloped property at any time during the life of the project. In doing so the developer collects a salvage value that is directly proportional to the current market value of the remaining potential development minus any remaining remediation costs. This proportion represents the part of the price for developed property that may be attributed to the value of the land given its entitlements. This particular property has been appraised at \$47,300,000. Given the definition of land value in (2.1) and the expectation of the entitlements from the distribution in Table 2.2 it may be seen that the proportion of the developed value which may be attributed to the land is 0.16.

A discount rate of 8% has been selected which remains consistent with previous work in the area of real estate development. For reference the model parameters and their values are presented in Table 2.1 along with a list of variable descriptions.

2.4.2 Dynamics of Residential Real Estate Prices

Correctly specifying the dynamics of real estate prices is critical to obtaining an accurate valuation of the brownfield regeneration project. The traditional one factor Ito process considered in most real estate applications is limiting and tends to smooth out information about potential uncertainty. For example consider the standard geometric Brownian motion (GBM) process considered by Quigg [1995], Capozza and Li [2002], and Paxson [2007] among others. This process assumes that absent of uncertainty the market price will grow exponentially over time. With the addition of the diffusion component the process becomes stochastic allowing for the price to deviate from the trend. However, this assumption of GBM dynamics may be inadequate for describing housing prices that traditionally exhibit short term deviations from the expected long-run growth rate. Such movements are commonly described as housing "bubbles" or "booms" representing a short term increase in prices above trend, typically followed by a sharp decrease in the value [Case and Shiller, 2003]. In order to capture such dynamics we consider the two factor stochastic process $S_t = (P_t, \mu_t)^T$, where P_t follows a continuous time random walk with a stochastic mean reverting drift μ_t. The process represents the solution to the set of stochastic differential equations

$$dP_t = \mu_t P_t dt + \sigma_P P_t dW_1 \tag{2.20}$$

$$d\mu_t = \delta(\bar{\mu} - \mu_t)dt + \sigma_\mu dW_2, \tag{2.21}$$

where W_1 and W_2 are standard independent Brownian motions.[12]

In order to facilitate estimation of the parameters in the processes (2.20)-(2.21) we consider the affine multivariate process $\tilde{S}_t = (\ln(P_t), \mu_t)$, where

$$d\ln(P_t) = \left(\mu_t - \frac{1}{2}\sigma_P^2\right) dt + \sigma_P dW_1. \tag{2.22}$$

[12] It was initially assumed that the Brownian motions would be correlated, though there was no statistical evidence to suggest a significant correlation and therefore the the assumption of independence has been applied.

The affine nature of the process means that the conditional mean and variance are

$$E\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = \begin{bmatrix} \bar{\mu}\Delta + \frac{1}{\delta}\left(\bar{\mu} - \mu_t\right)\left(e^{-\delta\Delta} - 1\right) - \frac{1}{2}\sigma_p^2\Delta + \ln(P_t) \\ e^{-\delta\Delta}\left(\mu_t - \bar{\mu}\right) + \bar{\mu} \end{bmatrix} \tag{2.23}$$

and

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = \begin{bmatrix} \sigma_p^2\Delta + \frac{\sigma_\mu^2}{\delta^3}\left[e^{-\delta\Delta}\left(2 - \frac{1}{2}e^{-\delta\Delta}\right) + \frac{2\delta\Delta - 3}{2}\right] & \frac{\sigma_\mu^2}{2\delta^2}\left(e^{-\delta\Delta} - 1\right)^2 \\ \frac{\sigma_\mu^2}{2\delta^2}\left(e^{-\delta\Delta} - 1\right)^2 & \frac{\sigma_\mu^2}{2\delta}\left(1 - e^{-2\delta\Delta}\right) \end{bmatrix}. \tag{2.24}$$

We refer the reader to Appendix C for the derivation of the moments. The standard Kalman filter is used to compute the log-likelihood for the system using the normal conditional transition density and the mean and variance in (2.23)-(2.24). An optimization routine may then be applied in order to find the set of parameters which maximize the log-likelihood.

The parameters are estimated using the Office of Federal Housing Enterprise Oversight (OFHEO) housing price index for the Houston / Sugar Land / Baytown MSA from Q2 1976 to Q2 2007. We consider the data series only up to Q2 2007 as this was the point in time when our corporate partner was considering an investment in this particular project. There has been some debate as to the ability of indices maintained by the OFHEO to predict dynamics for the price of new development. The main argument against the use such indices is that they are created using preexisting home sales within urban centers, when the primary location for new housing development is at the urban fringe. This is not a concern within the context of this example as the brownfield in question lies in what can be considered an urban center. The index is inflation adjusted to 2007 U.S. dollars using the consumer price index from the Bureau of Labor Statistics. The estimates for the parameters of the multi-factor price process are presented in Table 2.3 with the standard errors in parenthesis.[13]

[13]The standard errors are computed using the diagonals of information matrix formed from the inverse of the Hessian at the optimum as described by Hamilton [1994].

Table 2.3: Parameter Estimates for the Housing Price Process

Parameter	**Estimate**
δ	0.2705 (0.0698)
$\bar{\mu}$	-0.0047 (0.0003)
σ_P	0.0415 (0.0000)
σ_μ	0.0231 (0.0002)

2.4.3 Numerical Results

Provided with the parameters of the empirical case study we compute a numerical approximation for the brownfield remediation and redevelopment project's value using the techniques detailed in Appendices A and B. The partial derivative of the value functions with respect to the non-stochastic state variables, that is the remaining construction in the development stage and the remediation remaining in the initial stage, are approximated using a forward finite difference on a discrete grid. The number of points on the grid is fixed at 100 for both approximations. The value at each of the discrete steps is approximated as a piecewise linear function with 400 breakpoints in the P direction and 100 in the direction of μ. The multidimensional function is represented as the tensor product of the univariate versions such that at each step of the deterministic variable, and in the case of the remediation stage given the state of the entitlement process, the value is defined by 40,000 approximating coefficients. Therefore, with eight possible outcomes for the regulation process and the resolution of the discrete grid for the forward finite differences, the value of the project is defined by over 36,000,000 unknown coefficients. The approximation covers the state space where $P \in [0, 300]$ and $\mu \in [-.25, .25]$.

Within the development stage the firm is considered to have to have control of a discrete variable dictating the current status of the investment, in addition to a continuous variable representing the relative price. The optimal management policy for these controls will be dependent upon the state as determined by the market price, P_t, the market growth rate, μ_t, and the development entitlements remaining, K_t. The optimal control policy for the relative price, θ^*, will therefore be a function over the three dimensional state space. Whereas the optimal policy with regards to the discrete control will be represented as a partition of the three-dimensional state space. In order to analyze these controls we examine cross sections of the optimal policies. We begin by considering the case where $K_t = \epsilon$, which would be the start of development stage when the firm receives all of the entitlements it requested. The optimal control policy for this situation is presented in Figure 2.2. In this plot the policy for the discrete choice variable defining the operating regime is defined by the switching boundaries. The policy for the continuous control defining the price markup is represented by a contour plot of the three dimensional function for θ^*. The choice to suspend construction may be seen as determined by the relationship between the growth rate of the real estate market and the firm's discount rate. When the growth rate variable is high, and in particular when it exceeds the discount rate, it is optimal for the firm to suspend operations and allow the market price to increase. Due to the mean reverting nature of the stochastic growth rate, it is expected to eventually decrease, at which point the firm will resume operations in the event that the price is high enough. In the case of low prices, it may be suboptimal for the firm to resume development even though the true growth rate of the market price is below the discount rate. Here the firm will remain suspended and monitor the market. If conditions improve enough, in terms of higher prices and possibly higher growth rates, the firm may restart operations. However, if market conditions do not improve but instead deteriorate, the firm will sell off the remainder of the project. Interpretation of the optimal policy for the continuous control θ is straightforward. As the growth rate decreases and becomes negative it is in the firm's best interest to decrease the sale to market price ratio in order to speed up the absorption of completed units, before the market loses too much value.

As the firm is actively developing, the remaining entitlements will be decreasing at a rate determined by the optimal sale price. While the state of K_t is declining the project's value will change to reflect this reduction in possible future construction. The developer's optimal control policy will also adjust. Figure 2.3 presents the evolution of the control policy as the firm completes the project, for the case where the growth rate variable is equal to 5 percent, $\mu_t = 0.05$. This value of μ is chosen because at this point in the state space there exists both a region for suspension and selling the project. At lower values of μ there only exists a trigger value for selling off the remaining undeveloped property, as eventually the the suspension region disappears as μ declines.

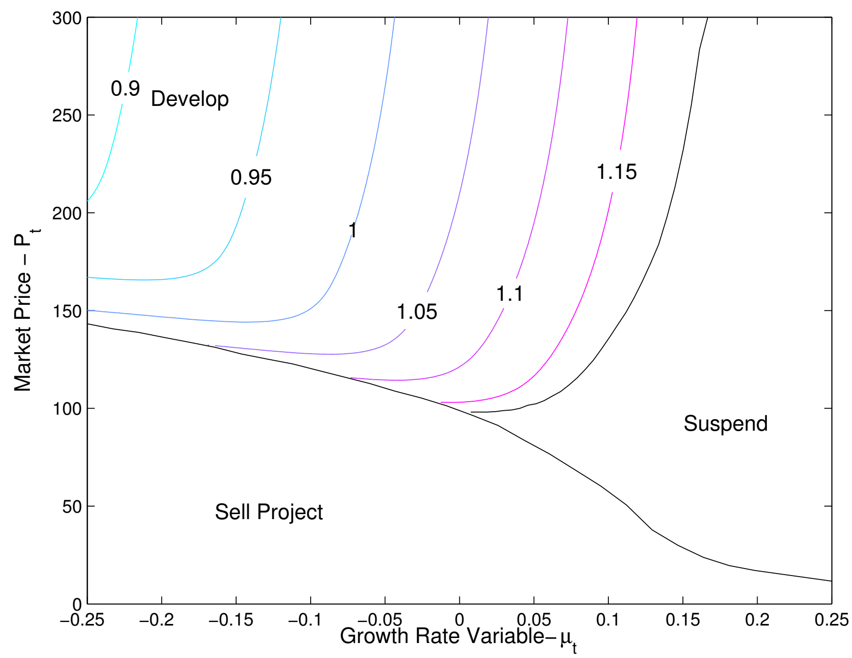

Figure 2.2: Development Stage Optimal Control Policy, K_t Fixed

Optimal control policy for the case where the remaining stock of entitlements equals the number requested, $K_t = \epsilon$. The policy for the discrete control variable is denoted by the switching boundaries and the policy for the continuous control in the relative price is representing by a contour plot.

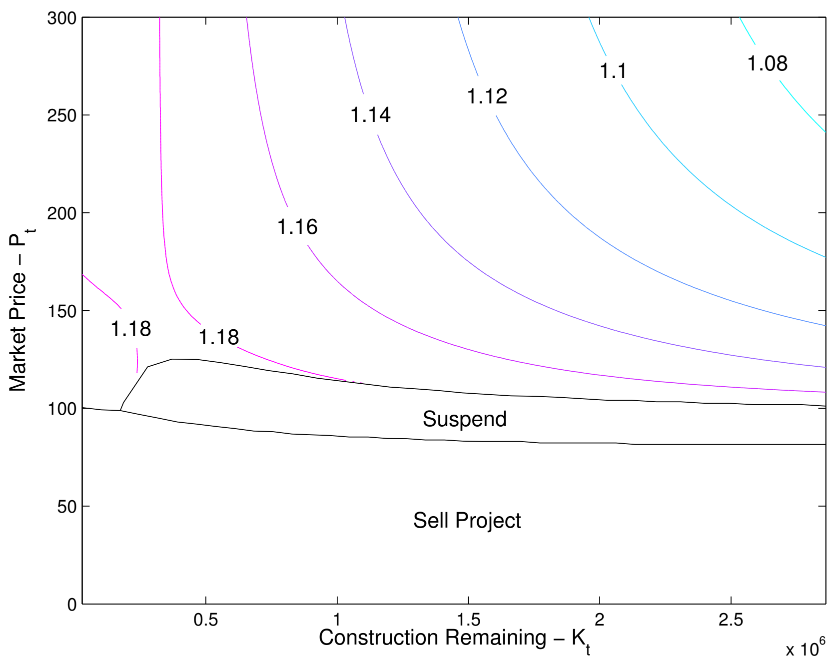

Figure 2.3: Development Stage Optimal Control Policy, μ Fixed

Optimal control policy for the case where the market price growth rate is $\mu_t = 0.05$. The policy for the discrete control variable is denoted by the switching boundaries and the policy for the continuous control in the relative price is representing by a contour plot.

As may be seen in Figure 2.3 the optimal policy for the relative price is generally as expected, increasing as the remaining units the firm may supply decreases. However when the project is near completion and the market price is relatively low, as compared to the marginal cost of construction, it becomes optimal to slightly reduce the sale price in order to increase the absorption rate. This behavior is an attempt to avoid suspending the project when the there are only a small number of entitlements remaining. This is because the costs associated with suspending the project begin to outweigh the benefits due to the limited number of units the firm would be able to build once they resume operations. This is why there exists a second contour for the relative price of 1.18. This behavior may be seen in the suspension boundary as well. Near the end of the project, that is when almost all entitlements have been used, we see that it is no longer as valuable to suspend the investment as is represented by the decrease in the trigger price. In fact as the project gets very close to completion it will never be optimal for the firm to suspend construction, even with the absence of switching costs. The benefit associated with being able to sell the few remaining units at a slightly higher price will not warrant the negative cash flow incurred while suspended. Similarly, the benefit to the firm of suspending operations at relatively low prices relative to the marginal cost of construction is also decreasing as the remaining development entitlements decrease. Therefore, the trigger price for selling the remainder of the project is increasing as K_t decreases.

Based on the optimal policy for the discrete and continuous control variables, a cross section of the value for the development stage is presented in Figure 2.4, for the case where the growth rate is equal to its long-run mean, $\mu_t = \bar{\mu}$. The points within the K_t space used for the cross section correspond to the possible outcomes of the entitlement process. In other words, given $\mu_t = \bar{\mu}$ and the market price, P_t, this represents the value of the option to develop which the firm will receive upon completion of the remediation and entitlement stage. The pronounced kink in the curves corresponds to the trigger price for selling the project as presented in Figure 2.3. The development value of the remediated and entitled site represents the driving force behind the value of the project in the initial stage and is independent of any uncertainty in the regulation process. The value function at alternative values of μ will have a similar shape, however the kink will be at a higher or lower price

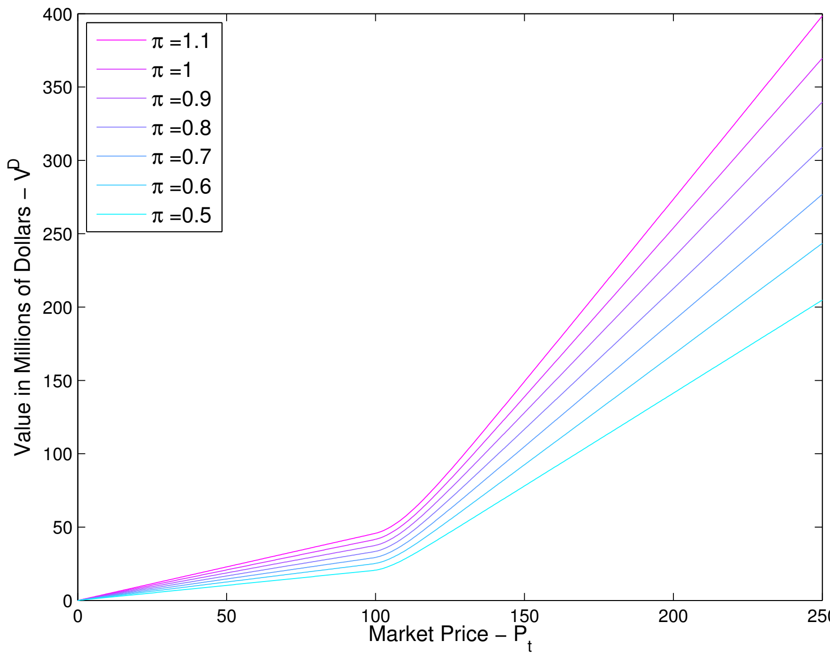

Figure 2.4: Development Value -V^D

Development stage value at the possible outcomes of the entitlement process given that the market growth rate is at its long-run mean, $\mu_t = \bar{\mu}$.

depending upon the optimal switching boundaries when μ is below or above its long-run mean.

Provided with the value of the development stage it is possible to approximate both the value and optimal control policy for the remediation and entitlement stage. The initial stage of the investment is defined in regards to four state variables representing the market price, market growth rate, remediation remaining, and the status of the entitlement process. In order to simplify the discussion we restrict ourselves to two dimensions by focusing on the beginning of the project. This corresponds to the state in which remediation has not yet begun, $R_t = \bar{R}$, and the regulation process has not yet produced an outcome, $Y_t = 0$. The optimal control policy for this state is represented by the switching boundaries in Figure 2.5. These boundaries differ from those in the development stage in an intuitive manner. Due to the lengthly nature of the remediation stage, it is no longer optimal to suspend operations simply due to the presence of a high growth rate. Instead it is now optimal for

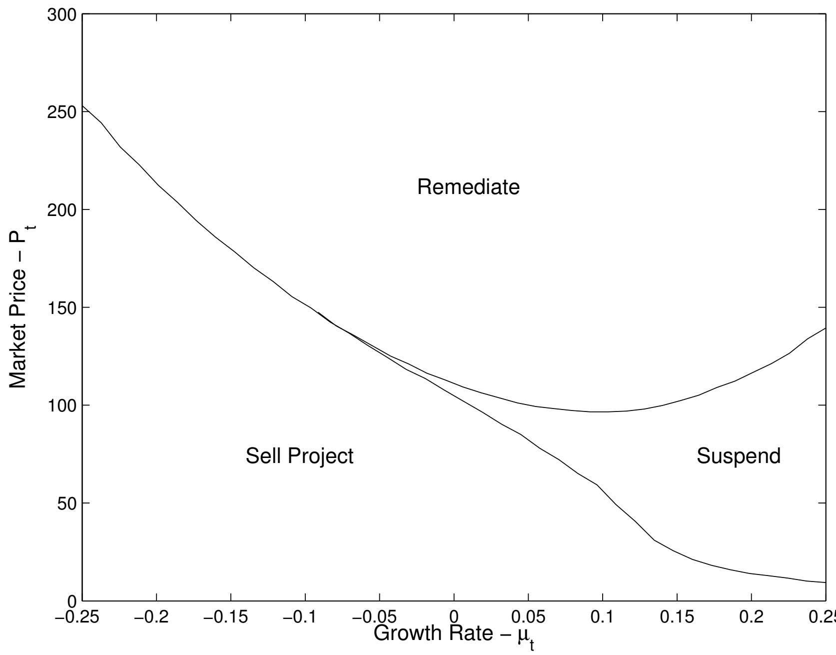

Figure 2.5: Remediation Stage Optimal Control Policy

Optimal control policy at the beginning of the project when all remediation remains, $R_t = \bar{R}$, and entitlements have not been received, $Y_t = 0$.

the firm to continue remediating the site while the market price rises from the high positive growth rate. If when the firm enters the development stage the growth rate variable μ is still relatively high compared to the discount rate, the firm may choose to suspend operations at that time. Additionally, due to the significant amount of time needed for remediation and the negative cash flow associated with the process, the boundary for selling the project now resides at a set of points corresponding to better market conditions than in the development stage. This is particularly evident in the case of of extremely low values of μ. Even at high prices, the decrease in development value over the life of the remediation stage as a result of the negative growth rate will make it optimal to sell the project. Based on the optimal control policy for the firm the value of the project at the beginning of its life is presented in Figure 2.6.

An important contribution of this paper is the inclusion of entitlement regulation within the brownfield regeneration process. Ignoring the risk associated with this part

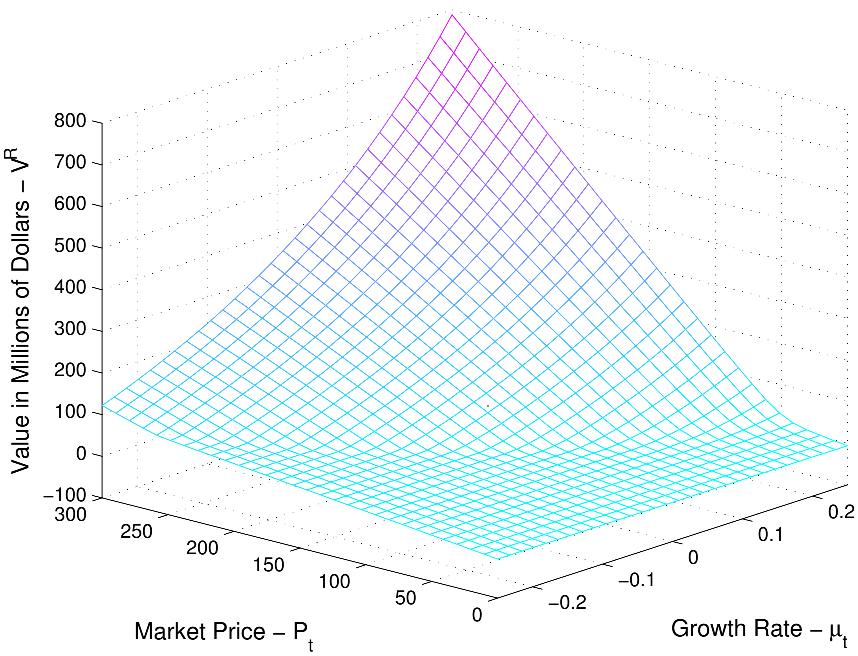

Figure 2.6: Remediation Stage Value - V^R

Value at the beginning of the project when all remediation remains, $R_t = \bar{R}$, and entitlements have not been received, $Y_t = 0$.

of the redevelopment process will lead to significant errors in valuing such investments. Excluding the process and assuming that all entitlements requested would be granted prior to development, that is $\pi = 1$ with a probability of one, would result in a significant overvaluation of the project. At the outset of this particular project it was estimated that $P_t = 150$ and $\mu = 0.0114$. At this point in the state space ignoring the entitlement process would result in an overvaluation of over 40 million dollars or 26% of the base case value. This extreme error is a result of ignoring the uncertainty associated with the outcome and length, and in turn cost of the regulatory procedure. For the example considered in this paper, the distribution governing the outcome of the regulatory process has significant downside risk as is evident from the expectation lying well below $\pi = 1$. Having local governments downsize the purposed development is common in these types of projects. Altering the assumption such that the developer will be able to construct units equivalent to the expected entitlements, $K_t = E[\pi]\epsilon$, will still result in a significant valuation error of 1.07 million dollars or 1.1% of the base case value. This approach will partially include the downside risk associated with the regulatory outcome, though it ignores the variance of the distribution. Figure 2.7 presents the overvaluation that would occur from ignoring the risk associated with the entitlement process and assuming the expected entitlements with no delay.

The variance of the distribution governing the regulatory outcome is an important property of the investment, as it has a significant impact of the project's value. An understanding of this relationship is imperative as the outcome's variability will differ markedly across regions. According to our corporate partner, a common scenario in Europe and certain communities in the United States (e.g. California) is one where the entitlement outcome has two possibilities. Approval for the project as designed, with the potential for small changes that may be considered negligible so that $\pi = 1$, or complete refusal for rezoning and permitting such that $\pi = 0$. This distribution for the regulatory outcome may be seen as an extreme in terms of the outcome's variance relative to the base distribution presented in Table 2.2. In order to maintain the same expected value the new distribution is set such that the outcomes $\pi = 0$ and $\pi = 1$ are associated with the probabilities 0.176 and 0.824 respectively. With this increased uncertainty in the outcome of the entitlement

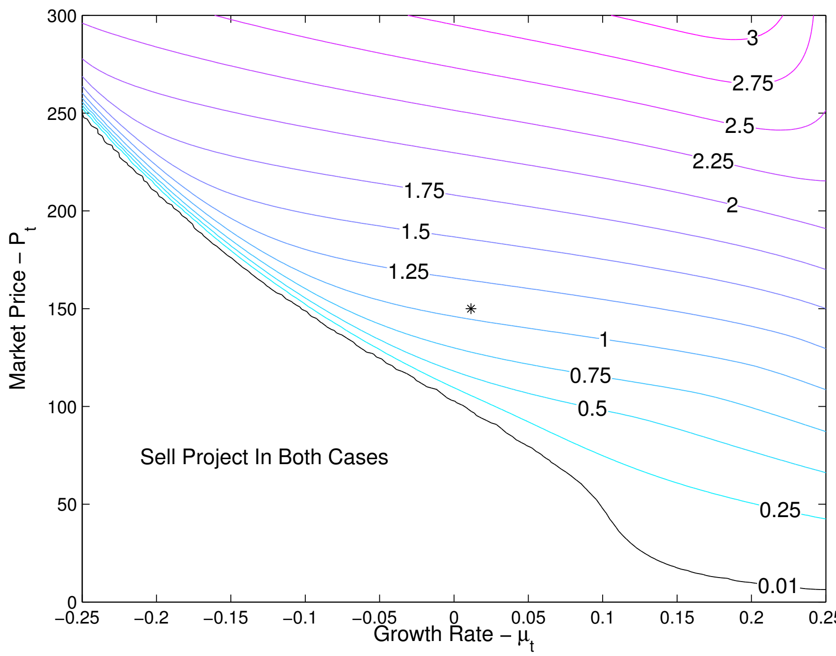

Figure 2.7: Overvaluation in Millions Assuming Expected Entitlements and No Delay

Given that all remediation remains, $R_t = \bar{R}$, the contours represent the difference in value (in millions of dollars) between the base case with the outcome distribution presented in Table 2.2 and $\lambda = 0.5$ and the case with no entitlement uncertainty or regulatory delay. The * represents the point where $P_t = 150$ and $\mu = 0.0114$, which was estimated to be the current point in the state space when the investment was being considered by our corporate partner in Q3 2007.

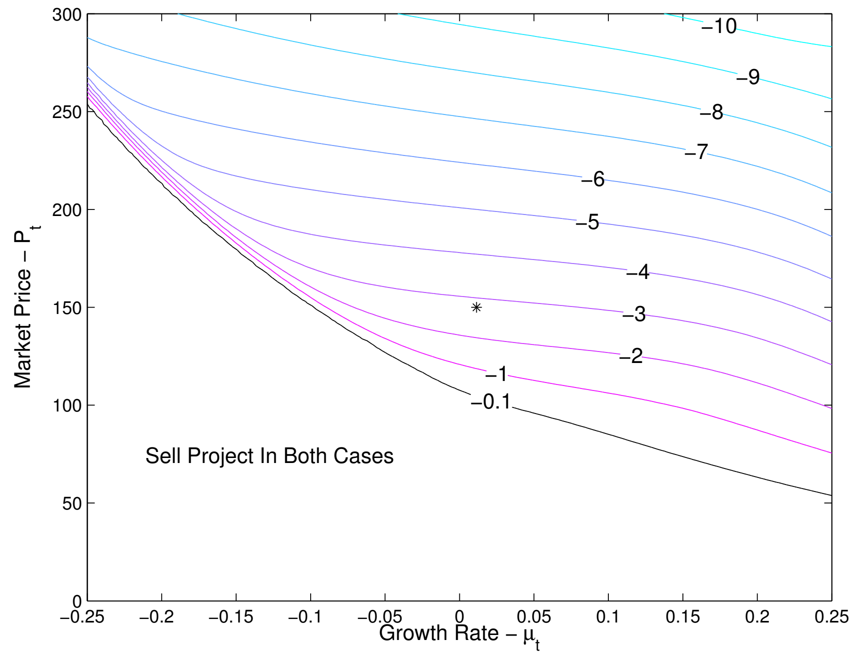

Figure 2.8: Value Change in Millions with All or Nothing Entitlements

Given that all remediation remains, $R_t = \bar{R}$, the contours represent the difference in value (in millions of dollars) between the base case with the outcome distribution presented in Table 2.2 and the "all or nothing" case with $\pi \in 0, 1$ with $Pr(\pi = 0) = 0.176$ and $Pr(\pi = 1) = 0.824$. The * represents the point where $P_t = 150$ and $\mu = 0.0114$, which was estimated to be the current point in the state space when the investment was being considered by our corporate partner in Q3 2007.

process the investment's value is significantly decreased in states where it is optimal for the firm to maintain ownership of the property under the base distribution. For the particular case of $P_t = 150$ and $\mu = 0.0114$ the increased uncertainty reduces the brownfields value by over 2.8 million dollars or 2.81% percent of the base case value. The impact of the increased variability across the entire state space is presented in Figure 2.8. This suggests that sites within communities associated with "all or nothing" approval practices have a disadvantage in attracting potential investors.

The other key parameter in the entitlement process is the expected length until development approval. Previous work with greenfield development has found that the regulation lag has a significant impact on the decision of developers to engage in new housing construction [Mayer and Somerville, 2000a]. This effect may be attributed to the impact that a lengthy entitlement process will have on the value of such real estate investments. In the case of brownfields the firm must also undertake the time consuming task of remediation prior to the development phase. It is therefore the relationship between the length of the processes that will determine the impact of the regulation. Since the actual time until the firm receives development entitlements is a random variable, we examine the effect of the expected length. For the base case the expected length of the regulatory process is 6 months, while the remediation process requires 2 years to complete. When $\mu = 0.0114$, Figure 2.9 presents the reduction in the project's value when the expected length of the entitlement process is extended to 1, 2, and 3 years. The shape of these curves is related to the optimal control policy for the firm. At low levels of the price it is optimal for the firm to sell off the project independent of the regulation process' expected length, and therefore it has no impact on the value. When the price is within the range where it is optimal for the firm to remain suspended the expected lag length will have little effect on the value outside the increased cost associated with the longer regulation procedure. However, when market conditions are such that the firm will currently be actively remediating, or would do so with a slight price increase, an escalation in the expected time until approval may have a significant impact on the project's value. For the case where the expected time until approval is increased from 6 months to 1 year the impact is relatively low, as this lag is still well below the 2 years needed for remediation. It is worth noting that this decrease in the value is still greater than the expected increase in costs associated with the longer entitlement procedure due to the uncertainty in the process. As the expected length is increased the probability of the process taking longer than the time needed for remediation increases. This is especially significant when the market price is high enough that the firm is likely to complete the remediation process without suspending, in order to begin the development stage as soon as possible. When the expected length of the regulatory proceedings is increased to $\bar{R}$ and beyond, the probability of the process taking longer than remediation is high when the price is high and therefore the value is seen to be lower. These results further illustrate that the inclusion of the regulatory process is critical in accurately valuing redevelopment projects. The temporal nature of the process along with its uncertainty result in a significant impact

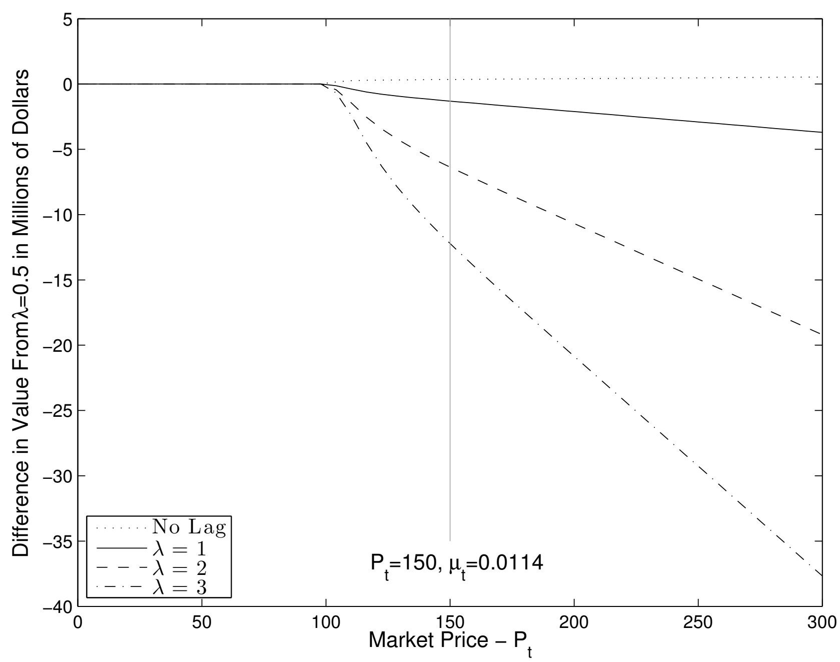

Figure 2.9: Effect of the Expected Length of the Entitlement Process

For the case where all remediation remains, $R_t = \bar{R}$, and $\mu = 0.0114$ the lines represent the difference in value (in millions of dollars) between the base case with $\lambda = 0.5$ and the various alternative values of λ. The vertical line represents the point where $P_t = 150$ and $\mu = 0.0114$, which was estimated to be the current point in the state space when the investment was being considered by our corporate partner in Q3 2007.

on the investment's value.

Again we consider the particular state of the market surrounding this project at its inception, where the price was \$150 and μ was estimated to be 0.0114. This case is represented in Figure 2.9 by the vertical line. As may be seen, the impact of increasing the expected lag from 6 months to 1 year is relatively small though when the length of the entitlement process becomes long compared to the time required for remediation the impact may be substantial. These results may also be seen in the fact that the increase in value from an expected length of 6 months to immediate notification of development entitlements is relatively small. To be specific the increase in the value, \$346,497, is only slightly greater than the expected cost of undergoing the 6 month regulatory procedure, \$250,000. The loss in value beyond the expected negative cash flow from the lengthier regulatory procedure may be seen as a result of the uncertainty in the process. Though, for the case where the expected lag goes from 6 months to 3 years the value decreases by \$12,215,420 or 12.4% of the base case value.

2.5 Concluding Remarks

This paper introduces the concept of a real options framework to value multistage investments where movement between stages requires both endogenous actions in addition to exogenous regulatory approval. The model is applied to the problem of valuing brownfield remediation and redevelopment projects, in which obtaining the right to develop the site requires not only the removal of environmental contamination but also entitlement approval. Within this regulatory process both the outcome and the time until approval exist as factors of uncertainty to the firm. We find through our case study that the expected length of the regulatory process does not have a significant impact on brownfield value due to the lengthy remediation process required to remove the environmental contamination. Though the presence of uncertainty in the time required to obtain entitlements will have a negative effect on the site's value since it introduces a positive, though very minimal, probability that the regulatory process will cause a delay in development. These results demonstrate that traditional methods used to encourage site development such as "fast-track" approval may not provide significant incentives for developers in the case of brownfields. Furthermore, this example has broader implications which suggest that investment valuation models used to assess the impact of public sector uncertainty, should be careful to consider the temporal dynamics of both regulatory and investment procedures.

On the other hand we find evidence that the uncertainty associated with the outcome of the entitlement process does in fact have a significant impact on the value of brownfields. Therefore when comparing potential investments remediation firms will choose sites within municipalities where they perceive there to be a lower variance in the regulatory outcome, ceteris paribus. This suggests that communities with the reputation for greater uncertainty in terms of approved entitlements, particularly those perceived to have "all or nothing" approval processes, may be required to offer higher incentive packages to attract investors. As seen with the case presented in Section 2.4 the impact of having a reputation for "all or nothing" approval may be substantial.

This paper provides a significant update to the literature on valuing investments in brownfields, not only through the introduction of the entitlement process but also with the inclusion of more realistic dynamics for real estate markets. The traditional assumption that prices follow GBM implies that the expected growth rate of the market price will be constant over time, and therefore decisions may be made based solely on the current price level. However, this result does not take into account developer expectations about short-term movements in the real estate market that will fluctuate over time due to uncertainty in the growth rate, a characteristic which should be considered by firms when making decisions. Therefore we introduce a two-factor process for real estate price and growth dynamics which captures such short term fluctuations. In turn our real options framework provides a description of firm behavior which is defined over a state space that is more representative of the market information that should be of concern to developers. As such it provides a more attractive framework for the analysis of features such as the entitlement process or incentive packages used to encourage redevelopment. The case study presented suggests that ignoring the role of short term growth rates on developer's decisions may lead to serious errors in valuing brownfield regeneration investments.

With the current abundance of brownfields, the complexity of such investments, and the multitude of benefits stemming from their redevelopment there exists both a need and an opportunity for more research within the area of valuing brownfield remediation projects. Firm's operating on medium to large sites typically incorporate both residential and commercial real estate development into such projects. This feature provides the potential opportunity for municipalities to increase the attractiveness of particular sites by offering flexible entitlements. Such entitlements allow the developer to alter the proportion of residential and commercial real estate during the development stage. This type of incentive for brownfield redevelopment is relatively new and its value has yet to be studied. The real options model for redevelopment presented in this paper provides ideal framework for the analysis of flexible entitlements in addition to other incentives.

Chapter 3

Solving Optimal Switching or Impulse Control Models Combined with a Stochastic Control Using Newton's Method

The use of stochastic control modeling in economics and finance has grown substantially over the last twenty years. The optimal switching framework has been of particular importance due to its ability to describe the behavior of economic agents under uncertainty. In this capacity it has been applied to a variety of problems ranging from firm investment patterns, to optimal resource extraction, to the management of invasive species.[1] In its basic form the optimal switching framework considers an economic agent who is dynamically choosing from a discrete set of operating environments, in an attempt to maximize a particular objective function. The agent moves between the regimes in response to changes in the state of nature, where the evolution is at least in part stochastic. We denote this type of control over the current regime as a discrete control because it may only take on one of countably many values.

[1] We refer the interested reader to Dixit and Pindyck [1994], Trigeorgis [1993b], and Schwartz and Trigeorgis [2004] for collections of applications utilizing the optimal switching and other stochastic control frameworks.

In certain cases however, it is natural to view the agent as having not only this type of discrete control, but also one or more continuous controls that may take on one of an infinite number of values. An excellent example of this class of hybrid control models is the application considered in the seminal paper of Brennan and Schwartz [1985]. They analyze the case of optimal resource extraction in the form of a firm operating a mine under an uncertain output price. The firm has a discrete control over the status of operations, whether the mine is open or closed, in addition to continuous control in the rate of extraction. This type of combined optimal switching and stochastic control model for natural resource management was extended by Lumley and Zervos [2001]. Other uses for such hybrid controls include the behavior of firms in imperfect markets. Chapter 2 examines the case of investments in brownfield remediation and redevelopment projects. During the development phase the firm has the ability to adjust the sale price through a continuous stochastic control but also holds a discrete control representing the options to suspend or abandon the project.

A significant issue with the implementation of the combined optimal switching and stochastic control framework in such applications is that the existence of the two control types is almost enough to ensure that a closed form solution to the problem will not exist. To date we unaware of any generalized numerical methods for solving this class of combined optimal switching and stochastic control problems. Typical are the simplifications utilized in Brennan and Schwartz [1985] and Lumley and Zervos [2001] where the continuous control variable is either fixed or discretized, thereby returning the problem to one in which there only exists the discrete control over the current operating regime. Such simplifications have previously allowed for approximate solutions to such hybrid control models, though at the cost of potentially serious errors in the project value and optimal control policies. In this paper we present a new numerical method to solve the general class of combined optimal switching and stochastic control problems. We rely on previous work to show that for this class of hybrid stochastic optimal control models the value function is a viscosity solution to the Hamilton-Jacobi-Bellman quasi-variational inequality, where the presence of the continuous control variable implies that the inequalities are non-linear second order partial differential equations with respect to the state variables. A numerical method is introduced that redefines the problem of approximating the value function and optimal control policies as an extended vertical non-linear complementarity problem (EVNCP) that may be solved using Newton's method. We demonstrate this technique using the optimal resource extraction application of Brennan and Schwartz [1985].

Closely related to the case of the combined optimal switching and stochastic control model is that of the combined impulse and stochastic control model. Applications of this framework include Oksendal and Sulem [2002] who consider the task of optimal portfolio management in the presence of fixed transaction costs. They demonstrate that an economic agent maximizing utility through control of consumption and their portfolio may be modeled as a combined impulse control and stochastic control model. Similar portfolio management applications have been considered by Zakamouline [2006], Chang [2007], and Ly Vath et al. [2007]. The combined impulse and stochastic control framework has also been applied to the problem of exchange rate management by Mundaca and Øksendal [1998] and Cadenillas and Zapatero [2000], where the government is considered to have a continuous stochastic control in the interest rate and an impulse control allowing them to intervene in the exchange market through the purchase and sale of currency.

These combined impulse and stochastic control models have primarily been solved using a distinct two step policy iteration procedure. The policy for the continuous control is fixed and the value function and optimal impulse control are then evaluated; these are then fixed and used to update the policy for the continuous control. This procedure is iterated until there is convergence in the optimal control policy. More recently, Øksendal and Sulem [2005] have demonstrated that the solution to a combined impulse and stochastic control problem satisfies a Hamilton-Jacobi-Bellman quasi-variational inequality and as such numerical methods should allow for the solution of the value function and control policies in one step. We extend this work by showing that the quasi-variational inequality for such problems may be redefined as an EVNCP similar to the approach taken in solving the combined optimal switching and stochastic control case. The portfolio management example of Oksendal and Sulem [2002] is used as an example to demonstrate the technique. An alternative approach to solving combined impulse and stochastic control problems is to redefine the problem as combined optimal switching and stochastic control problem using the method of Balikcioglu [2008]. This method is demonstrated using the exchange rate control application of Cadenillas and Zapatero [2000].

The paper proceeds as follows: Section 3.1 defines the combined optimal switching and stochastic control model and presents the main numerical method; Section 3.2 presents an extension to the case of combined impulse and stochastic control models; Section 3.3 presents evidence of the procedure's merits using examples from the literature; and Section 3.4 provides some concluding remarks.

3.1 Combined Stochastic Control and Optimal Switching

We consider an optimal switching model in which the principal agent will be operating in one of m regimes. The current regime in time t is denoted by Z_t where

$$Z_t \in \{z_1, \ldots, z_m\} =: \mathcal{Z}.$$

The choice of the current regime will be denoted as a discrete control held by the agent as it may only take on one of countably many values.

Along with the choice of regime the agent is assumed to have control over an additional n continuous variables denoted by $X_t \in \mathbb{R}^n$ where S represents the set of state variable. This continuous control will have an affect on the drift of the stochastic process S_t, which describes the state of the environment for the agent. It is assumed that $S_t \in \mathbb{R}^d$ is the solution to the time homogeneous[2] stochastic differential equation

$$dS_t = \mu(S_t, Z_t, X_t)dt + \sigma(S_t, Z_t)dW, \tag{3.1}$$

where $\mu : \mathbb{R}^d \times \mathcal{Z} \times \mathbb{R}^n \to \mathbb{R}^d$ is the drift function and $\sigma : \mathbb{R}^d \times \mathcal{Z} \to \mathbb{R}^{d \times d}$ is the diffusion function. To ensure the existence of a unique solution to (3.1) it is assumed that $\mu(\cdot, z, x)$

[2] The analysis in this paper assumes a time homogenous stochastic process for the state variables. However, the model may be extended by defining one of the state variables as time with a constant drift of one and a diffusion of zero.

and $\sigma(\cdot, z)$ are Lipschitz continuous and that

$$|\mu(s, z, x)|^2 + ||\sigma(s, z)||^2 \leq D(1 + |s|^2)$$

for some constant D, where $|\cdot|$ and $||\cdot||$ are the vector and matrix norm respectively, therefore ruling out explosive growth [Oksendal and Karsten, 1998]. $W_t = W(t, \omega); t \geq 0, \omega \in \Omega$ is a d-dimensional Brownian motion in the $(\Omega, \mathcal{F}, \mathbb{P})$ probability space with the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. In this paper we do explicitly incorporate the possibility of general Lévy processes, though note that the method may be easily extended to handle the case of jump diffusions.

There exist p constraints on the continuous control, X_t, such that

$$g(S_t, Z_t, X_t) \leq 0. \tag{3.2}$$

The control choice x is said to be admissible if (3.2) is satisfied, where the admissible set is denoted as $\mathcal{A}(S_t, Z_t)$ such that for a given state, $S_t = s$, and regime, $Z_t = z$,

$$x \in \mathcal{A}(s, z) = \{x : g(s, z, x) \leq 0\}. \tag{3.3}$$

A policy w for the discrete control variable may be described as the possibly finite double sequence

$$w = (\tau_1, \tau_2, \ldots, \tau_k, \ldots; \zeta_1, \zeta_2, \ldots, \zeta_k, \ldots),$$

where $0 \leq \tau_1 \leq \tau_2 \leq \ldots$ are stopping times with respect to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. Associated with the stopping times are the changes to the discrete control $\zeta_k \in \mathcal{Z}$. In other words at time τ_k the agent switches to the regime ζ_k. Therefore a policy for the combined continuous and discrete control may be written as $\nu = (w, x)$. The combined policy ν is considered to be admissible if (3.3) holds, $\zeta_k \in \mathcal{Z}\ \forall k$, and τ_k are stopping times, where the set of all admissible combined controls is denoted as V.

The complete state of the system may be described by the stochastic process

$Y_t \in \mathbb{R}^d \times \mathcal{Z}$ such that

$$Y_t = \begin{pmatrix} S_t \\ Z_t \end{pmatrix}.$$

When an admissible control $\nu \in V$ is applied, the system has the form

$$Y_t = Y_t^\nu = \begin{pmatrix} S_t^\nu \\ \zeta_k \end{pmatrix} \qquad \text{if } \tau_k \leq t < \tau_{k+1},$$

where S_t^ν represents the dynamics for the states variables in S_t given the control policy ν.

When the system is in the state $y = (s, z)$, with the continuous control policy $x \in \mathcal{A}(s, z)$, the agent receives a flow of benefits $f(s, z, x)$ where $f : \mathbb{R}^d \times \mathcal{Z} \times \mathbb{R}^n \to \mathbb{R}$. Given the state s the cost of switching from regime z to ζ is denoted by $C(s, z, \zeta)$ where $C : \mathbb{R}^d \times \mathcal{Z} \times \mathcal{Z} \to \mathbb{R}$. For completeness note that there is no cost to remaining in the current regime, hence $C(s, z, z) = 0$. The expected discounted flow of benefits to the agent at time zero under the policy ν may then be defined as

$$J^\nu(s, z) = E\left[\int_0^\infty e^{-rt} f(S_t^\nu, Z_t, X_t)dt - \sum_{j=1}^{\infty} e^{-r\tau_j} C(S_{\tau_j}, \zeta_{j-1}, \zeta_j) \middle| S_0 = s, Z_0 = z)\right],$$

where $Z_t = \zeta_j$ if $\tau_j \leq t < \tau_{j+1}$, $\zeta_0 = z$, and r is the discount rate. The problem is then to find for all $y = (s, z)$ the value function $V(s, z)$ such that

$$V(s, z) = \sup_{\nu \in V} J^\nu(s, z). \tag{3.4}$$

Given the existence of an optimal combined control policy $\nu^* = (w^*, x^*)$ the value function may be defined as

$$V(s, z) = J^{\nu^*}(s, z).$$

As a natural extension of the work by Øksendal and Sulem [2005] on combined impulse and stochastic control models and Brekke and Oksendal [1994] on switching models, the conditions that define the value function, $V(s, z)$, for the combined optimal switching and stochastic control model are

$$V(s,z) \geq \sup_{\zeta \in \mathcal{Z}\setminus\{z\}} \left[V(s,\zeta) - C(s,z,\zeta)\right] \tag{3.5}$$

and

$$rV(s,z) \geq \sup_{x \in \mathcal{A}(s,z)} \left[f(s,z,x) + \mathcal{L}^{z,x}V(s,z)\right], \tag{3.6}$$

where $\mathcal{L}^{z,x}$ is the differential generator

$$\mathcal{L}^{z,x} = \sum_{i=1}^{d} \mu_i(s,z,x)\frac{\partial}{\partial s_i} + \frac{1}{2}\sum_{i=1}^{d}\sum_{j=1}^{d}\left[\sigma(s,z)\sigma(s,z)^T\right]_{ij}\frac{\partial^2}{\partial s_i \partial s_j}, \tag{3.7}$$

and either (3.5) or (3.6) must hold with equality. If (3.5) is the condition to hold with equality then it will be optimal for the agent to switch regimes. The regime to which they move will be determined by the supremum. On the other hand if (3.6) holds with equality, then it is optimal to remain in the current regime. These conditions represent a combination of the Hamilton-Jacobi-Bellman (HJB) equation for stochastic controls and the quasi-variational inequality (QVI) for optimal switching problems, and therefore we denote (3.5)-(3.6) as a HJBQVI following the similar work of Øksendal and Sulem [2005].

We assume that a solution for the optimal control policy x^* exists within the admissible set $\mathcal{A}(s,z)$. This solution will then satisfy the Karush-Kuhn-Tucker necessary conditions

$$\nabla_x f(s,z,x^*) + \sum_{i=1}^{d} \nabla_x \mu_j(s,z,x^*)\frac{\partial V(s,z)}{\partial s_j} + \sum_{i=1}^{n} \lambda_i \nabla_x\, g_i(s,z,x^*) = 0, \tag{3.8}$$

$$g_i(s,z,x^*) \leq 0 \quad \forall i = 1,...,p, \tag{3.9}$$

$$\lambda_i \geq 0 \quad \forall i = 1,...,p, \tag{3.10}$$

and

$$\lambda_i g_i(s,z,x^*) = 0 \quad \forall i = 1,...,p. \tag{3.11}$$

It may therefore be noted that the optimal continuous control, x^*, is a function of the state in addition to partial derivatives of the value function such that

$$x^* = x^*\left(s, z, \frac{\partial V(s,z)}{\partial s}\right). \tag{3.12}$$

Given the notation for the optimal continuous control in (3.12) the HJBQVI may be rewritten as

$$V(s,z) \geq \sup_{\zeta \in \mathcal{Z} \setminus \{z\}} [V(s,\zeta) - C(s,z,\zeta)] \tag{3.13}$$

and

$$rV(s,z) \geq f(s,z,x^*) + \mathcal{L}^{z,x^*} V(s,z). \tag{3.14}$$

As before one of the conditions in (3.13)-(3.14) must hold with equality, which one determines the optimal policy for the discrete control, while the optimal policy for the continuous control will be defined by (3.12).

3.1.1 Numerical Solution

A closed form solution for the value function defined by (3.13)-(3.14) does not generally exist though accurate numerical approximations may be obtained for most problems through the use of projection methods. The value function is approximated as

$$V(s,z) \approx \phi(s)c_z,$$

where ϕ represents a set of q basis functions for a family of approximating functions, and c_z is a q-dimensional vector of coefficients for the value associated with the z^{th} regime. Given this approximation the condition in (3.14) may be rewritten as

$$r\phi(s)c_z - f(s,z,x^*) - \mathcal{L}^{z,x^*}\phi(s)c_z \geq 0, \tag{3.15}$$

Note that condition (3.15) is not linear in the coefficient vector c_z due to the fact that the optimal control policy definition in (3.12) includes the first derivative of the value function such that

$$x^* = x^* \left(s, z, \frac{\partial \phi}{\partial s_1} c_z, \ldots, \frac{\partial \phi}{\partial s_d} c_z \right).$$

The additional quasi-variational inequality condition in (3.13) may be rewritten as

$$\min_{\zeta \neq z} \left[\phi(s)(c_z - c_\zeta) - C(s, z, \zeta) \right] \geq 0. \tag{3.16}$$

Since one of the conditions in (3.15) and (3.16) must hold with equality the problem may be restated as

$$\min \left(r\phi(s)c_z - f(s, z, x^*) - \mathcal{L}^{z,x^*}\phi(s)c_z, \min_{\zeta \neq z} \left[\phi(s)(c_z - c_\zeta) - C(s, z, \zeta) \right] \right) = 0. \tag{3.17}$$

The goal here is to solve for the $q \times m$ unknown values of the coefficient vectors c_z $\forall z \in \mathcal{Z}$.

Collocation Solution

The problem of finding the coefficients that satisfy (3.17) for each of the m regimes given a set of q nodal points may be described as an extended vertical non-linear complementarity problem (EVNCP) of the form

$$\min \left(F_1(\tilde{c}), F_2(\tilde{c}), ..., F_m(\tilde{c}) \right) = 0, \tag{3.18}$$

where the mq-dimensional vector $\tilde{c}$ represents the coefficient vectors, c_z $\forall z \in \mathcal{Z}$, stacked vertically. Let Φ represent the set of basis functions evaluated at q nodal points, denoted by s. Then the mq-dimensional functions $F_i(\tilde{c})$ derived from (3.17) are given by

$$F_i(\tilde{c}) = e_i \otimes H_i(\tilde{c}) + \left[(\mathrm{I}_m - \underline{1}_m e_i^T) \otimes \Phi \right] \tilde{c} + \begin{pmatrix} C(s, 1, i) \\ \cdots \\ C(s, m, i) \end{pmatrix}, \tag{3.19}$$

where $\underline{1}_m$ is a column vector of m ones and e_i is the i^{th} column of the $m \times m$ identity matrix, I_m, and

$$H_i(\tilde{c}) = r\Phi c_i - f\left(s, i, \hat{x}\right) - L^{i,\hat{x}} c_i,$$

where the matrix $L^{z,x}$ is the differential generator $\mathcal{L}^{z,x}$ evaluated at the q nodal values s and[3]

$$\hat{x} = x^* \left(s, i, \frac{\partial \phi}{\partial s_1} c_i \bigg|_s, \ldots, \frac{\partial \phi}{\partial s_d} c_i \bigg|_s \right). \tag{3.20}$$

The min function in (3.18) represents a piecewise linear function that is not continuously differentiable, and therefore regular Newton-type methods may not be able to determine the appropriate descent direction [Sun and Qi, 1999]. Therefore we can redefine the EVNCP using a semi-smooth function with the same roots as the min function. This may be accomplished through the application of an NCP-function $\gamma : \mathbb{R}^2 \to \mathbb{R}$ where

$$\gamma(a, b) = 0 \Longleftrightarrow a, b \geq 0, \quad ab = 0.$$

The EVNCP in (3.18) requires an extension to $\Gamma : \mathbb{R}^m \to \mathbb{R}$ in order to accommodate the m regimes in the model. This may be accomplished by iterating the NCP-function such that

$$\Gamma_i = \gamma(\Gamma_{i-1}, d_i),$$

where $\Gamma_1 = d_1$ and $\Gamma_m = \Gamma(d_1, d_2, ..., d_m)$.[4] A well studied choice for the NCP-function is

$$\gamma(a, b) = a + b - \sqrt{a^2 + b^2}, \tag{3.21}$$

introduced by Fischer [1992] and commonly denoted as the Fischer-Burmeister function. Using the iterated NCP-function the EVNCP in (3.18) is reformulated as

$$\Gamma[F_1(\tilde{c}), F_2(\tilde{c}), ..., F_m(\tilde{c})] = 0. \tag{3.22}$$

The system of nonlinear equations in (3.22) may then be used to compute the coefficients of the value function approximation through the use of a Newton based algorithm. Given the semi-smooth property of the NCP-function, the Newton iteration will converge to the solution $\tilde{c}^*$, where $\Gamma[F_1(\tilde{c}^*), F_2(\tilde{c}^*), ..., F_m(\tilde{c}^*)] \approx 0$ [De Luca et al., 1996].

---

[3] We note that when finite difference methods are used to compute the derivatives of the basis functions that $L^{z,x}$ is has a slightly different interpretation than $\mathcal{L}^{z,x}$ evaluated at the q nodal values.

[4] The iteration may be performed in any order.

3.1.2 Computing the Jacobian

The solution method laid out in Section 3.1.1 may be rewritten to allow for a more computationally efficient implementation. Consider rewriting (3.19) as

$$F_i(\tilde{c}) = M_i\tilde{c} + l_i - U_i(\tilde{c}), \tag{3.23}$$

where

$$M_i = e_i e_i^T \otimes \left[r\Phi - \frac{1}{2}\sum_{i=1}^{d}\sum_{j=1}^{d}\left[\sigma(s,z)\sigma(s,z)^T\right]_{ij} \left.\frac{\partial^2\phi}{\partial s_i \partial s_j}\right|_s \right] + (\mathrm{I}_m - \underline{1}_m e_i^T) \otimes \Phi, \tag{3.24}$$

$$l_i = \begin{pmatrix} C(s,1,i) \\ \cdots \\ C(s,m,i) \end{pmatrix}, \tag{3.25}$$

and

$$U_i(\tilde{c}) = e_i \otimes \left\{ f\left(s,i,\hat{x}\right) + \sum_{i=1}^{d} \mu_i(s,z,\hat{x})\frac{\partial\phi(s)}{\partial s_i}c_i \right\}, \tag{3.26}$$

where $\partial\phi(s)/\partial s_i$ is a q-vector evaluated at the nodal points. By decomposing the function into its linear and nonlinear components it is possible to reduce the computation time by pre-computing M_i and l_i for every regime. Therefore during each iteration of the generalized Newton method only the nonlinear component U_i must be recomputed.

This decomposition will also allow for a more efficient way to compute the Jacobians of the EVNCP arguments, $F_i(\tilde{c})$. Based on the definition in (3.23) it may be seen that

$$\frac{\partial F_i(\tilde{c})}{\partial\tilde{c}} = M_i + \frac{\partial U_i(\tilde{c})}{\partial\tilde{c}}. \tag{3.27}$$

Therefore in each iteration the computation of the Jacobian only requires computing the derivative of the nonlinear component. As follows from the definition in (3.26)

$$\frac{\partial U_i(\tilde{c})}{\partial\tilde{c}} = e_i \otimes \left\{ \left[\frac{\partial f(s,i,\hat{x})}{\partial x} + \sum_{j=1}^{d} \frac{\partial\mu_i(s,i,\hat{x})}{\partial x} \left.\frac{\partial\phi}{\partial s_i}\right|_s c_i \right] \frac{\partial\hat{x}}{\partial\tilde{c}} + \sum_{j=1}^{d}\mu_i(s,i,\hat{x}) \left.\frac{\partial\phi}{\partial s_i}\right|_s \right\}. \tag{3.28}$$

The first component on the right hand side will be equal to zero since for an interior solution the first order condition in (3.8) will hold, and $\partial\hat{x}/\tilde{c} = 0$ otherwise. Therefore (3.28) reduces to

$$\frac{\partial U_i(\tilde{c})}{\partial \tilde{c}} = e_i \otimes \left[\sum_{j=1}^{d} \mu_i(s,i,\hat{x}) \left. \frac{\partial \phi}{\partial s_i} \right|_s \right].$$

In addition to implementing the decomposition of $F_i(\tilde{c})$ one may increase the performance of the algorithm by noting that M_i and $\partial U_i(\tilde{c})/\partial\tilde{c}$ are relatively sparse. This is particularly true when standard finite difference methods are applied to approximate the value function. By exploiting the sparsity of these components one is able to reduce the memory requirements and reduce the computational burden associated with the linear solve for determining the step direction in the generalized Newton algorithm.

3.1.3 Relationship to Policy Iteration

It is worth noting that the Newton based method introduced in this paper is equivalent to a policy iteration algorithm, such as the one purposed by Øksendal and Sulem [2005] to solve the class of combined impulse and stochastic control models. To show this relationship we consider the k^{th} step for the policy iteration algorithm. First the policy for both the continuous and discrete variables are updated and then used to derive a linear system $G^k\hat{c}^k + l^k = 0$ that defines the updated approximating coefficients for the value function. The optimal policy for the continuous control variable is updated using (3.20) such that

$$x^k = x^* \left(s, z, \left. \frac{\partial \phi}{\partial s_1} c_z^{k-1} \right|_s, \ldots, \left. \frac{\partial \phi}{\partial s_d} c_z^{k-1} \right|_s \right) \quad \forall z \in \mathcal{Z}. \tag{3.29}$$

Subsequently the policy for the discrete control is updated using the HJBQVI in (3.18) where

$$\min \Big(F_1(\hat{c}^{k-1}), F_2(\hat{c}^{k-1}), \ldots, F_m(\hat{c}^{k-1}) \Big) = 0,$$

where $\hat{c}^{k-1}$ represents the the m approximation coefficient vectors from step $k-1$ stacked vertically and $F_i(\cdot)$ is as defined in (3.19). If the i^{th} component of minimization problem is the smallest for row j, then it will be the case that

$$(G^k)_{j\cdot} = (e_i \otimes (r\Phi - \Delta^{i,x^k}) + [(I_m - \underline{1}_m e_i^T) \otimes \Phi])_{j\cdot}, \tag{3.30}$$

and

$$(l^k)_j = \left(-e_i \otimes f(s,i,x^k) + \begin{pmatrix} C(s,1,i) \\ \cdots \\ C(s,m,i) \end{pmatrix} \right)_j . \tag{3.31}$$

The approximating coefficients are then updated using the linear system $G^k \hat{c}^k + l^k = 0$ such that

$$\hat{c}^k = -(G^k)^{-1} l^k. \tag{3.32}$$

This process is repeated until there is convergence in the value function.

Now we consider the k^{th} step of Newton's method when using the min function instead of the semi-smooth version derived from the Fischer-Burmeister function. Again the first step is to compute the continuous control policy using (3.29), which is then substituted into the HJBQVI. Then the residual from the HJBQVI is computed where the j^{th} element is equal to $(G^k)_{j\cdot}\hat{c}^{k-1} + (l^k)_j$, where G^k and l^k are defined as in (3.30) and (3.31) when the i^{th} element of the HJBQVI is the smallest. The Jacobian of the HJBQVI will be equal to G^k for the reasons discussed in Section 3.1.2, and therefore the approximating coefficients for the k^{th} Newton iteration will be

$$\begin{aligned} \hat{c}^k &= \hat{c}^{k-1} - (G^k)^{-1}(G^k \hat{c}^{k-1} + l^k) \\ &= -(G^k)^{-1} l^k. \end{aligned}$$

This represents the same update that is performed in policy iteration algorithm, in (3.32).

3.2 Combined Stochastic Control and Impulse Control

The same numerical technique laid out in Section 3.1 may be extended to the case of combined stochastic control and impulse control models. In both cases a viscosity solution is defined by a HJBQVI for which an approximate solution may be obtained by solving an

EVNCP. There are however, distinct differences in the components of the EVNCPs due to the differences between a switching and impulse control. While the previous case of combined optimal switching and stochastic controls has not previously been addressed, the problem of solving combined impulse and stochastic control problems has received some attention. Most notable is the work by Øksendal and Sulem [2005], which defines the HJBQVI for the combined control problem and suggests the use of finite difference methods to approximate the value function. In this section we add to this line of research by introducing the specifics of a numerical method for this class of hybrid control models. In particular we demonstrate how numerical techniques similar to those presented in Section 3.1 may be used to define the problem of approximating a solution as an EVNCP that may be solved using Newton's method. We begin with a description of the framework and the HJBQVI defining the value function and optimal control policy, followed by the setup of the EVNCP for the combined impulse and stochastic control problem.

The state of the system will be defined by the stochastic process $S_t \in \mathbb{R}^d$ whose dynamics in the absence of an impulse control are described by

$$dS_t = \mu(S_t, X_t)dt + \sigma(S_t)dW, \tag{3.33}$$

where $\mu : \mathbb{R}^d \times \mathbb{R}^n \to \mathbb{R}^d$ is the drift function and $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d}$ is the diffusion function. Again, to ensure the existence of a unique solution to (3.33) it is assumed that $\mu(\cdot, x)$ and $\sigma(\cdot)$ are Lipschitz continuous and that

$$|\mu(s, x)|^2 + ||\sigma(s)||^2 \leq D(1 + |s|^2)$$

for some constant D, where $|\cdot|$ and $||\cdot||$ are the vector and matrix norm respectively, therefore ruling out explosive growth [Oksendal and Karsten, 1998]. $W_t = W(t, \omega)$; $t \geq 0, \omega \in \Omega$ is a d-dimensional Brownian Motion. The continuous control is represented by X_t and again may be subject to a series of constraints There exist n constraints on the continuous control, X_t, such that

$$g(S_t, X_t) \leq 0. \tag{3.34}$$

The set of admissible controls $\mathcal{A}(s, z) \subseteq \mathbb{R}^n$ is therefore defined by the conditions in (3.34).

In addition to the stochastic control the agent is assumed to have an additional control over the state of the system, where at any time τ the agent is able to exert a control over the state S_τ in the form of an impulse $\xi_\tau \in \mathcal{E} \subseteq \mathbb{R}^p$ such that

$$S_\tau = S_\tau^- + P(\xi_\tau),$$

where $S_\tau^- = \lim_{t \uparrow \tau} S_t$ and $P(\xi_\tau) : \mathcal{E} \to \mathbb{R}^d$. The set of admissible impulses $\mathcal{E} = \mathcal{E}(s)$ ensures that any constraints on the state of the system hold. When the agent exerts an impulse control ξ given the state s they pay a cost of $K(s, \xi)$ where $K : \mathbb{R}^d \times \mathcal{E} \to \mathbb{R}$. A policy w for the impulse control may be described as the possibly finite double sequence

$$w = (\tau_1, \tau_2, \ldots, \tau_k, \ldots; \xi_1, \xi_2, \ldots, \xi_k, \ldots),$$

where $0 \leq \tau_1 \leq \tau_2 \leq \ldots$ are stopping times with respect to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$. Associated with the stopping times are the impulse controls $\xi_k \in \mathcal{E}$. In other words at time τ_k the agent exerts the control ξ_k on the state of the system. Therefore a policy for the combined stochastic and impulse control may be written as $\nu = (x, w)$. The combined policy ν is considered to be admissible if (3.34) holds and $\xi_i \in \mathcal{E}(S_{\tau_i})$ $\forall i$. The set of all admissible combined controls is denoted as V.

When the system is in the state s, with the continuous control policy $x \in \mathcal{A}(s, z)$, the agent receives a flow of benefits $f(s, x)$ where $f : \mathbb{R}^d \times \mathcal{A}(s, z) \to \mathbb{R}$. Given the discount rate r the expected discounted flow of benefits to the agent at time zero under policy ν may be defined as

$$J^\nu(s) = E\left[\int_0^\infty e^{-rt} f(S_t^\nu, x_t)dt - \sum_{j=1}^\infty e^{-r\tau_j} K(S_{\tau_j}^-, \xi_j) \,\middle|\, S_0 = s\right].$$

The objective of the agent is find the optimal control policy that will maximize the discounted flow of benefits, such that

$$V(s) = \sup_{\nu \in V} J^{\nu}(s),$$

where $V(s)$ is denoted as the value function. Given the existence of the optimal control policy $\nu^* = (x^*, w^*)$ the value function may be defined as

$$V(s) = J^{\nu^*}(s).$$

The solution to this optimization problem is defined by the HJBQVI

$$\mathcal{V}^x(s) \geq \mathcal{M}\mathcal{V}^x(s) \tag{3.35}$$

and

$$r\mathcal{V}^x(s) \geq \sup_{x \in \mathcal{A}(s,z)} \left[f(s,x) + \mathcal{L}^x \mathcal{V}^x(s) \right], \tag{3.36}$$

where $\mathcal{M}$ is the intervention operator

$$\mathcal{M}h(s) = \sup_{\xi \in \mathcal{E} \setminus \{0\}} h\left[s + P(\xi)\right] - K(s, \xi), \tag{3.37}$$

and $\mathcal{L}^x$ is the differential generator for the combined stochastic and impulse control problem

$$\mathcal{L}^x = \sum_{i=1}^{d} \mu_i(s,x) \frac{\partial}{\partial s_i} + \frac{1}{2} \sum_{i=1}^{d} \sum_{j=1}^{d} \left[\sigma(s)\sigma(s)^T\right]_{ij} \frac{\partial^2}{\partial s_i \partial s_j}. \tag{3.38}$$

Where one of the conditions in (3.35)-(3.36) must hold with equality. Which condition holds with equality will again determine the optimal management policy for the agent. If (3.36) holds with equality then it is optimal for the agent to refrain from exerting an impulse control. However, if (3.35) holds with equality then the agent should exert an impulse control, whose value will be determined by the intervention operator. Øksendal and Sulem [2005] provides a proof that under mild conditions the value function is a viscosity solution to the HJBQVI in (3.35)-(3.36).

We assume that a solution for the optimal control policy x^* exists and is within the admissible set $\mathcal{A}(s,z)$. This solution will then satisfy the Karush-Kuhn-Tucker necessary conditions such that

$$\bigtriangledown_x f(s,x^*) + \sum_{j=1}^{d} \bigtriangledown_x \mu_i(s,x^*) \frac{\partial V(s)}{\partial s_j} + \sum_{i=1}^{n} \lambda_i \bigtriangledown_x g_i(s,x^*) = 0, \tag{3.39}$$

$$g_i(s,x^*) \leq 0 \quad \forall i = 1,...,n, \tag{3.40}$$

$$\lambda_i \geq 0 \quad \forall i = 1,...,n, \tag{3.41}$$

and

$$\lambda_i g_i(s,x^*) = 0 \quad \forall i = 1,...,n. \tag{3.42}$$

Once again it may be seen that the optimal control policy x^* will be a function of not only the current state but also the first partial derivatives of the value function, such that

$$x^* = x^*\left(s, \frac{\partial V(s)}{\partial s}\right). \tag{3.43}$$

The HJBQVI for the value function may then be rewritten as

$$V(s) \geq \mathcal{M}V(s) \tag{3.44}$$

and

$$rV(s) \geq f(s,x^*) + \mathcal{L}^{x^*}V(s), \tag{3.45}$$

where one of the conditions must hold with equality, $\mathcal{M}$ is the intervention operator in (3.37) and $\mathcal{L}^x$ is the differential generator in (3.38). It should be noted that dependence of optimal control policy, x^*, on the value functions first partial derivatives means that the differential generator in (3.45) will be a non-linear second order partial differential equation of degree d.

3.2.1 Numerical Solution

In order to solve for the value function and the optimal control policy in the combined stochastic and impulse control model we employ an approach similar to that presented in Section 3.1.1. Again the goal is to represent the value function in such a manner that the quasi-variational inequalities may be restated as a EVNCP that may be smoothed and solved using a Newton type method. Therefore we begin by approximating the value function using a set of q basis functions from a family of approximating functions, denoted as ϕ where

$$V(s) \approx \phi(s)c.$$

The q dimensional vector c represents the approximation coefficients. Given this approximation for the value function the condition in (3.45) may be rewritten as

$$r\phi(s)c - f(s, x^*) - \mathcal{L}^{x^*}\phi(s)c \geq 0. \tag{3.46}$$

The optimal policy for the continuous control variable x^* is defined by (3.43) using the approximating function such that

$$x^* = x^*\left(s, \frac{\partial \phi}{\partial s_1}c, \ldots, \frac{\partial \phi}{\partial s_d}c\right).$$

The additional condition in (3.44) may be restated as

$$\phi(s)c - \mathcal{M}\phi(s)c \geq 0. \tag{3.47}$$

The approximation of the quasi-variational inequality in (3.46)-(3.47) may be rewritten as

$$\min\left(r\phi(s)c - f(s, x^*) - \mathcal{L}^{x^*}\phi(s)c, \phi(s)c - \mathcal{M}\phi(s)c\right) = 0.$$

The goal is therefore to find the unknown vector of approximating coefficients c. This may be accomplished by restating the problem as a EVNCP where the approximation is evaluated at a set of q nodal points s. The complementarity problem is then defined as

$$\min\left(r\Phi c - f(s, x^*) - L^{x^*}c, \Phi c - \mathcal{M}\Phi c\right) = 0, \tag{3.48}$$

where Φ represent the set of basis functions evaluated at the q nodal points and the interpretation of L^{x^*} is the same as in Section 3.1. The EVNCP in (3.48) may be restated as a semi-smooth function using the Fischer-Burmeister function in (3.21). Newton's method may then be used to compute the vector of coefficients c.

Of particular concern when working with impulse control models is the intervention operator $\mathcal{M}$. The computational efficiency of this or any other algorithm that solves impulse control problems relies heavily on the efficient implementation of this optimization problem. While it is possible to implement a true optimization algorithm to solve the maximization problem associated with the intervention operator, this would be prohibitively costly since it would need to be performed for each nodal point during each evaluation of the smoothed EVNCP residual. An attractive alternative is to discretize the problem such that

$$\mathcal{M}\phi(s)c = \max_{\xi \in \mathcal{X}(s)} \phi[s + P(\xi)]c - K(s, \xi),$$

where $\mathcal{X}(s)$ is a discrete set of admissible interventions such that $\mathcal{X}(s) \subseteq \mathcal{E}(s) \backslash \{0\}$. Computing the intervention operator is then simply a search over the discrete set of impulse controls for the maximum.

A natural way to select the discrete set $\mathcal{X}(s)$ is to consider the set of impulse controls that would be admissible and would result in a new state that is also one of the nodal points used in the approximation. A serious problem with this approach is that there is no guarantee that the set of impulse controls associated with moving to another nodal points will be within the admissible set. This is the case with the example considered in Section 3.3.3, in which, for any nodal point on the discrete grid representing the state space, the set of new states that would result from admissible controls may be represented as a piecewise linear function with a discontinuity at the current state. For this specific example it may be the case that this function does not intersect with any points on the discrete grid, implying that $\mathcal{X}(s) = \{\}$, which is clearly not the case. This scenario is presented in Figure 3.1. Noting that this problem may extend beyond this specific example we purpose that in general the discrete set of admissible interventions, $\mathcal{X}(s)$, be explicitly selected as a discrete subset of the admissible controls $\mathcal{E}(s) \backslash \{0\}$.

As an alternative, Balikcioglu [2008] presents a method of representing a subclass

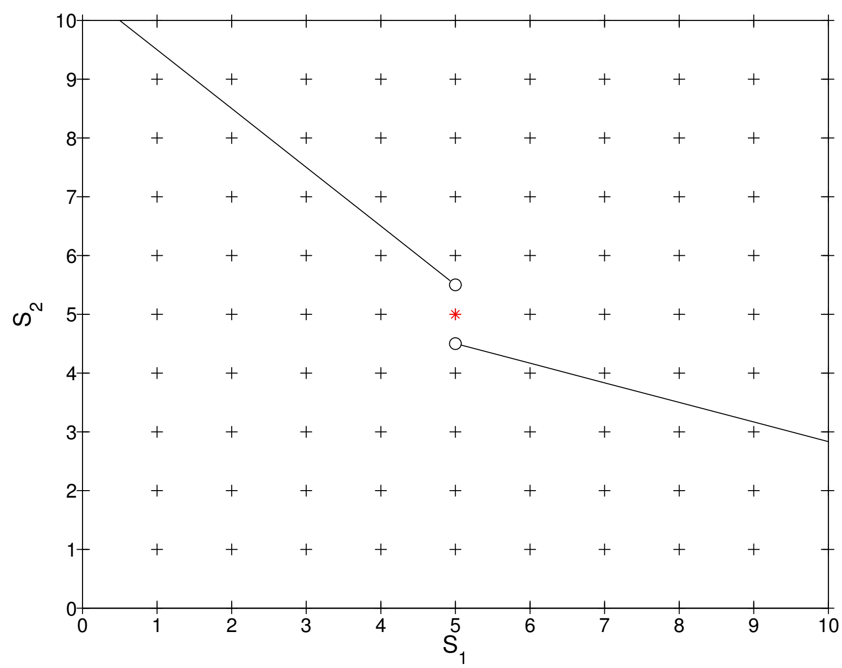

Figure 3.1: Admissible Set and Nodal Points

Depicts an example in which at the point $S = (5, 5)$, denoted by *, the admissible set, denoted by the solid line, does not intersect with any of the nodes on the discrete grid, denoted by $+$. This example is adapted from the problem considered in Section 3.3.3.

of impulse control problems as optimal switching problems, therefore avoiding the need to implement the intervention operator. In this case the problem of a combined stochastic and impulse control model would fall back to a combined stochastic control and optimal switching model which may be solved using the techniques from Section 3.1.

3.3 Examples

3.3.1 Optimal Resource Extraction

To examine these methods we consider the example of a firm operating a mine. This example is similar to those presented in Brennan and Schwartz [1985] and Brekke and Oksendal [1994]. The firm has control of a mine with Q_t units of ore remaining at time t and will operate the mine as to maximize the discounted flow of profits. When the ore is extracted it may be sold at the market price, P_t, which is assumed to follow geometric Brownian motion so that

$$dP_t = \alpha P_t dt + \beta P_t dW_t.$$

The firm is considered to have a continuous control in the choice of the proportional extraction rate, h_t, such that $h_t \in \mathbb{R}_+$. The cost of operating the mine at the current rate, $K(h_t)$, is assumed to be

$$K(h_t) = k_0 + k_1 h_t^2.$$

With $k_1 > 0$ this definition implies a standard upward sloping marginal cost curve. As may be noted, even at an extraction rate of zero, the firm will still incur some expenditures, denoted by $k_0 > 0$. These represent operating expenditures associated with keeping the mine open. The firm is assumed to have the option to forgo these expenses by closing the mine for a fixed cost of C_{21}. Once closed the mine may be reopened by the firm at a cost of C_{12}. These two states, open or closed, represent the two possible regimes in the model where

$$R = \begin{cases} 1 & \text{if the mine is closed} \\ 2 & \text{if the mine is open} \end{cases}.$$

Given this description of the problem the dynamics of the mine's ore may be described as

$$dQ_t = \begin{cases} 0 & \text{if } R = 1 \\ -h_t Q_t dt & \text{if } R = 2 \end{cases},$$

and the flow of rewards to the firm may be described as

$$f(P_t, h_t, R) = \begin{cases} 0 & \text{if } R = 1 \\ h_t Q_t P_t - K(h_t) & \text{if } R = 2 \end{cases}.$$

Once the firm chooses to open the mine they will operate at the optimal extraction rate, h^*. From the variational inequality in (3.13)-(3.14) it may be shown that the firm will suspend operation prior to choosing an extraction rate of zero if $k_0 - rC_{21} > 0$, where r is the discount rate. In other words, if the cost flow associated with remaining open and not extracting is greater than the opportunity cost associated with closing the mine it will never be optimal for the firm to set the extraction rate to zero. For the numerical example considered this condition is met and therefore the optimal extraction rate will be an interior solution such that

$$h^* = \frac{\left[P_t - V_Q(P_t, Q_t, R)\right] Q_t}{2k_1}, \tag{3.49}$$

where V_Q is the partial derivative of the value function with respect to the ore remaining. It should be noted that this holds only for the case where the firm is actively operating the mine, $R = 2$, as when the mine is closed, $R = 1$, the extraction rate is zero by definition.

The parameter values used in this example are presented in Table 3.1. The family of approximating functions used are piecewise linear with centered finite difference derivatives in the interior of the state space and one sided second order finite difference derivatives on the boundaries. We do not impose any constraints at these boundaries. We use 150 breakpoints for both state variable P and Q with a maximum value of \$20 for P and 100 units of ore for Q. For both state variables the minimum value considered is 0. The problem is then solved using a Newton algorithm with an Armijo line search for the step size, where the convergence criteria for the residual is set to 1^{-8} in the $\mathcal{L}^2$ norm.[5]

[5] The algorithm is based on the one detailed in Chapter 2 of Kelley [2003].

Table 3.1: Parameters for the Mine Example

Parameter	**Description**	**Value**
α	Market price drift rate	0.01
β	Market price diffusion rate	0.02
r	Discount rate	0.04
k_0	Fixed operating cost	3
k_1	Extraction cost parameter	200
C_{12}	Cost of opening the mine	5
C_{21}	Cost of closing the mine	5

A contour plot for the value of the mine in the open regime, $R = 2$, is presented in Figure 3.2. As anticipated the the value is increasing in both the market price and the amount of ore remaining in the mine. Figure 3.3 presents the optimal switching boundaries. The area of hysteresis between the two boundaries is due to the fixed cost of closing and opening the mine. Simply put, at a given level of ore it is optimal for the firm to wait until the price is far enough above the shutdown value before opening, in order to avoid paying the fixed cost when there exists a high probability of the price returning to below the shutdown value in the future.

When the mine is open the firm will extract the ore at the optimal rate as defined by (3.49). Figure 3.4 presents both the switching boundary and the contour lines for the optimal extraction rate. Below the switching boundary the optimal rate is zero due to the fact that in this area the firm would choose to close the mine.

An alternative approach to solving models with combined discrete and continuous controls, is to discretize the continuous control in order to utilize known methods of solving optimal switching models. In the simplest form one might consider a two regime model. In the first the mine is closed, and in the second the firm is extracting ore at a constant rate. In order to better capture the value of the firm's operating flexibility the practitioner may choose to include additional regimes that represent operating the mine at different extraction rates. For this example we consider three cases in which the continuous control has been discretized. The first case represents the two regime model where the firm is extracting ore at a rate of $h = 0.3$ when the mine is open. In the second case the firm is given greater flexibility through the use of a four regime model with extraction rate possibilities defined as $h \in \{0.2, 0.3, 0.4\}$. The third case provides for slightly greater flexibility with six regimes,

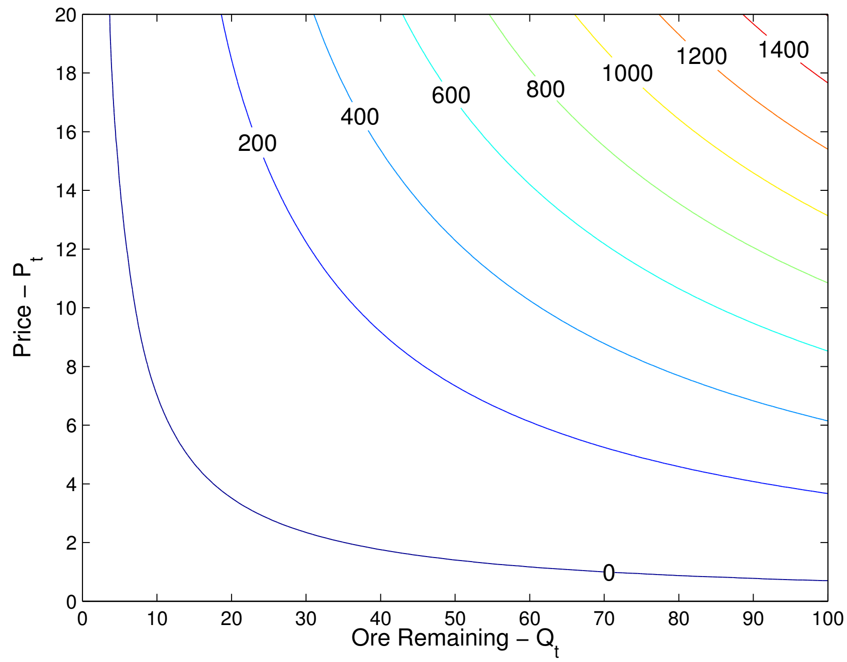

Figure 3.2: Value of the Mine

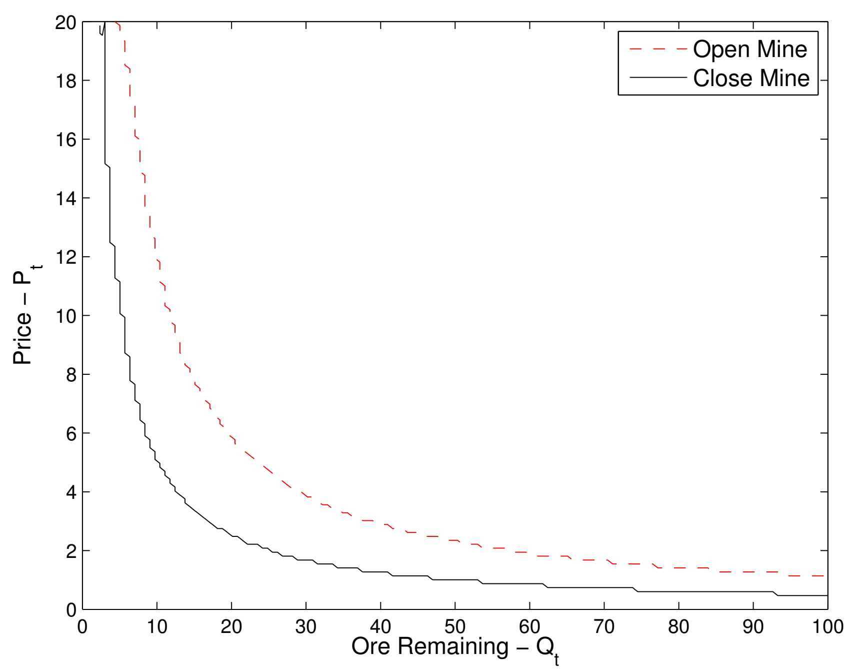

Figure 3.3: Switching Boundaries for Mine Operation

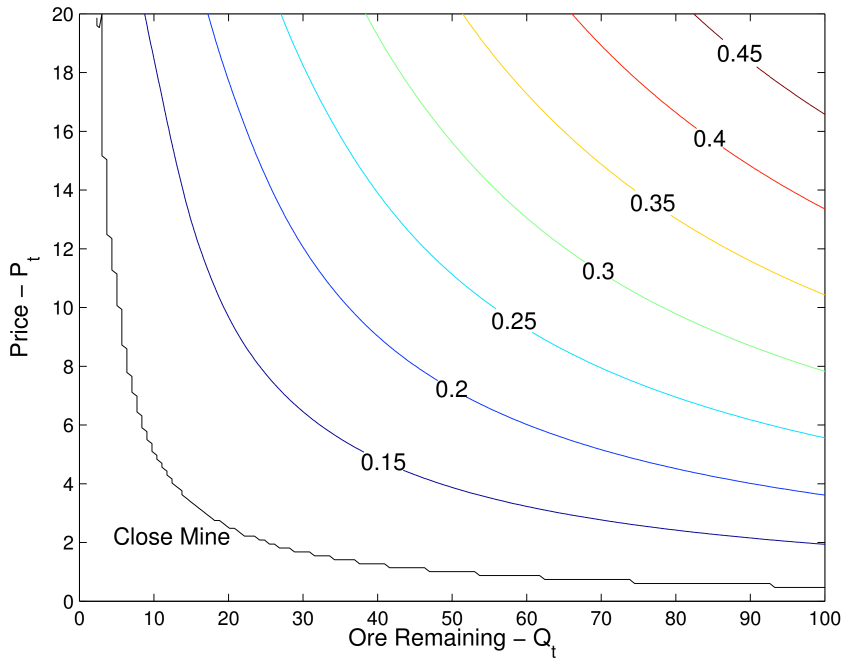

Figure 3.4: Optimal Control Policy While the Mine is Open

where $h \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$.[6] In all cases it is considered to be costless to move among the regimes in which the mine is actively operating. In other words there is no fixed cost associated with adjusting the extraction rate within the limited range of the control. As the model grows and allows for more flexibility in the choice of extraction rate the result will approach that of the case with the continuous control. However, as the number of regimes within the model increases so will the computational burden.

As the firm gains more flexibility with respect to the extraction rate the overall value of the mine will increase, reflecting the value of this control. Therefore, the models with the discretized control will always undervalue the project compared to the approach with the continuous control representing the case of full flexibility. To demonstrate this characteristic we evaluate the value function in all four cases using a grid of nodes with 100 evenly spaced points in both the P and Q state spaces respectively. We measure how closely

[6] We note that even though it is not optimal to chose an extraction rate of 0.5 within the [0,100]x[0,100] state space for the continuous control, it will be an optimal choice in the case of the discretized control for a subset of the state space.

the value of the discretized model with H choices of the extraction rate, $V^H(P,Q,R)$, approximates the value of the continuous control case, $V^C(P,Q,R)$ using the mean percentage difference over a discrete grid of 100 points in each direction of the state space, such that

$$\text{Mean Value Error} = \frac{1}{100}\sum_{k=1}^{100}\sum_{l=1}^{100}\frac{V^H(P_k,Q_l,2)-V^C(P_k,Q_l,2)}{V^C(P_k,Q_l,2)}.$$

In all cases the value is computed using the second regime. For the cases with more than one regime $V^H(P,Q,j)=V^H(P,Q,i) \quad \forall j,i \geq 2$ due to the fact that it is costless to switch between these regimes.

The results for this experiment are presented in Table 3.2.[7] It may be seen that the model with the continuous control takes only slightly longer to solve than the model with the fixed extraction rate. This is due to the fact that the existence of the continuous control requires very little additional computation per iteration as discussed in Section 3.1.2. One must only compute the optimal control policy for the extraction rate and subsequently update the non-linear component of the EVNCP. However, the more complex hybrid framework does take two additional Newton iteration, though this additional computational time appears minor compared to the the fact that the model with the fixed extraction rate fails completely in capturing the value of the firm's ability to adjust the extraction rate. The value of having flexibility in the extraction rate is evident by the high error when it is fixed. Through the addition of extra regimes to approximate the continuous control one is able to reduce the valuation error as seen by the decreasing norm. However, increasing the number of regimes also increase increase the dimensionality of the complementarity problem defining the solution which results in a large computational burden. It may clearly be seen that the additional time required by the non-linearity in the complementarity problem for the case of the continuous control is relatively small compared to the addition of more regimes. These results suggest that the numerical method presented in this paper will provide a more accurate solution for the combined optimal switching and stochastic control model than the case of a discretized continuous control. In addition solving the problem with the

[7]The runtime for the various models has been normalized to make comparison easier. We note that a value of 100 is equivalent to approximately 21 seconds on an Intel Core Duo 2.0 Ghz processor with 2 GB of RAM, running OS X 10.5.6 and MATLAB 2007a.

Table 3.2: Continuous vs. Discrete Control of the Extraction Rate

Model	**Iterations**	**Time**	**Mean Value Error**
Continuous Control	19	100.00	—
$h = 0.3$	17	75.33	-15.13%
$h \in \{0.2, 0.3, 0.4\}$	31	236.12	-3.82%
$h \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$	53	618.07	-0.67%

continuous control may be quicker than the case of the discretized control.

3.3.2 Exchange Rate Control

In this section we consider the application presented in Cadenillas and Zapatero [2000]. The model represents a central bank that seeks to maintain an exchange rate target through the use of an interest rate control and interventions in the foreign exchange market. This may be classified as a combined impulse and stochastic control problem, where the interest rate represents the continuous control and the market intervention is the impulse control. We transform this problem into a combined optimal switching and stochastic control problem and solve for the value function and optimal control policies using the method of Section 3.1.

The goal of the central bank is to maintain the target exchange rate ρ. The cost of deviating from the target is represented by the function

$$f(X_t, u_t) = (X_t - \rho)^2 + ku_t^2,$$

where X_t is the market exchange rate, u_t is the interest rate continuous control, and k is the marginal cost parameter for the interest rate control. The exchange rate, X_t, is assumed to follow a process that behaves like geometric Brownian motion when no controls are being exerted, such that

$$dX_t = (\mu X_t - Ku_t)\, dt + \sigma X_t dW_t.$$

In addition to the continuous control, the central bank may intervene in the foreign exchange markets directly in order to alter the exchange rate. Specifically, the central bank may exert an impulse ξ such that

$$X_\tau = X_\tau^- + \xi,$$

where X_τ^- is the value of the exchange rate just prior to the impulse. In the case where the central bank pushes the exchange rate upwards they face a fixed intervention cost of C and a marginal cost of c, such that the total cost of an impulse $\xi > 0$ is $C + c\xi$. Alternatively, if the impulse is designed to push the exchange rate downwards the fixed and marginal costs are denoted by D and d respectively.

The goal of the central bank is determine the optimal control policy that minimizes the expected discounted flow of costs associated with deviating from the target rate and implementing the control policies such that

$$V(X_t) = \sup_{u,\nu} E\left[\int_0^\infty -e^{-\lambda t} f(X_t, u_t)dt - \sum_{j=1}^{\infty} e^{-\lambda \tau_j} v(\xi_j)\right],$$

where λ is the discount rate, ν is the policy for the impulse control

$$\nu = (\tau_1, \tau_2, \ldots, \tau_n, \ldots; \xi_1, \xi_2, \ldots, \xi_n, \ldots),$$

and

$$v(\xi) = \begin{cases} C + c\xi & \text{if } \xi > 0 \\ D + d\xi & \text{if } \xi < 0 \end{cases}.$$

Following Section 3.2 the value function will satisfy the complementarity problem

$$\min\Big(\lambda V(X_t) - f(X_t, u^*) - \mathcal{L}^{u^*} V(X_t), V(X_t) - \mathcal{M}V(X_t)\Big) = 0,$$

where $\mathcal{L}$ is the differential generator and $\mathcal{M}$ is the intervention operator

$$\mathcal{M}V(X_t) = \sup_{\xi} V(X_t + \xi) - v(\xi),$$

and u^* is the optimal policy for the continuous control defined by

$$u^* = \frac{K}{2k} V'(X_t).$$

Extending the work of Balikcioglu [2008] to the combined impulse and stochastic control case, we transform this problem into a combined optimal switching and stochastic control problem with three regimes where

$$R = \begin{cases} 1 & \text{No impulse is being exerted} \\ 2 & \text{The rate is being pushed upward} \\ 3 & \text{The rate is being pushed downward} \end{cases}$$

where the value function, $V(X_t, R)$, satisfies the complementarity problem

$$\min\left(Q(X_t, u^*), V(X_t, 1) - V(X_t, 2) + C, V(X_t, 1) - V(X_t, 3) + D\right) = 0, \tag{3.50}$$

$$\min\big(V(X_t, 2) - V(X_t, 1), c - V'(X_t, 2)\big) = 0, \tag{3.51}$$

$$\min\big(V(X_t, 3) - V(X_t, 1), d + V'(X_t, 3),\big) = 0, \tag{3.52}$$

where $Q(X_t, u^*) = \lambda V(X_t, 1) - f(X_t, u^*) - \mathcal{L}V(X_t, 1)$. We refer the interested reader to the work of Balikcioglu [2008] for an in depth discussion of the intuition behind this transformation.

The combined optimal switching and stochastic control problem in (3.50)-(3.52) may easily be solved using the numerical method outlined in Section 3.1. In this example we use the same parameter values as in Cadenillas and Zapatero [2000]; these are presented in Table 3.3. The family of approximating functions used are piecewise linear with upwind finite difference derivatives and 1000 breakpoints. In the Newton algorithm the convergence criteria is set to 1^{-8}.

The approximate value functions are presented in Figure 3.5. It will be the case that at a trigger value of $b = 1.936$ it will be optimal for the central bank to exert a downward impulse to the target value of $\beta = 1.617$. For the upward impulse the trigger

Table 3.3: Parameters for the Exchange Rate Example

Parameter	Description	Value
μ	Uncontrolled exchange rate drift parameter	0.1
σ	Exchange rate diffusion parameter	0.3
λ	Discount rate	0.06
ρ	Exchange rate target	1.4
k	Marginal cost parameter for continuous control	1
K	Continuous control effectiveness	$-0.2\sqrt{2}$
C	Fixed cost of upward impulse	0.1672506173
c	Marginal cost of upward impulse	0.8
D	Fixed cost of downward impulse	0.015714985
d	Marginal cost of downward impulse	1.2

value is $a = 0.594$ with an associated target value of $\alpha = 0.937$. The optimal control for the continuous interest rate control is presented in Figure 3.6.

3.3.3 Portfolio Management

This example of a combined stochastic and impulse control model focuses on the portfolio management problem presented in Chancelier et al. [2002] and Oksendal and Sulem [2002]. An investor holds two assets, a risk-less investment X_t that grows at rate r and a risky investment Y_t whose dynamics are described by geometric Brownian motion such that

$$dY_t = \alpha Y_t dt + \sigma Y_t dW_t.$$

The investor may transfer the amount ξ from the risk-less asset to the risky asset for a fixed cost k and a proportional cost λ. Therefore the total cost of exerting the impulse ξ is $k + \lambda |\xi|$. The costs are taken from the account holding the risk-less asset so that exerting the impulse ξ at time τ will therefore affect the state such that

$$X_\tau = X_\tau^- - \xi - k - \lambda |\xi|$$

and

$$Y_\tau = Y_\tau^- + \xi.$$

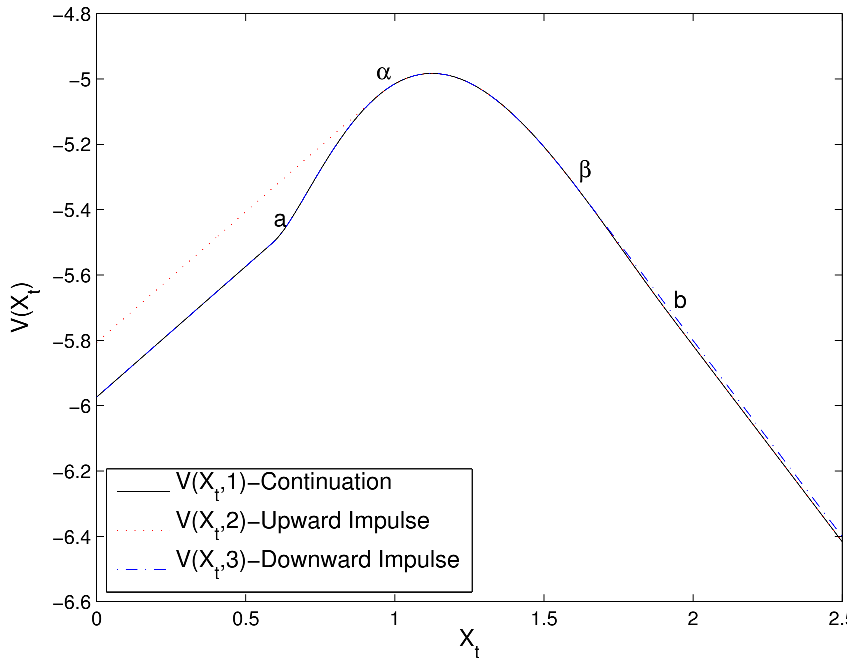

Figure 3.5: Value Functions for the Exchange Rate Example

A positive impulse $\xi > 0$ implies a purchase of the risky asset and a negative impulse $\xi < 0$ represents a sale of the risky asset. An admissible impulse ensures that the state variables must stay within the solvency region $X_t \geq 0$ and $Y_t \geq 0$, therefore prohibiting borrowing and selling short. The set of all admissible controls given the state (x, y) is denoted by $\Xi = \Xi(x, y)$. In terms of the notation from Section 3.2 the impulse control may be described by

$$P(\xi) = \begin{pmatrix} -k - \lambda\,|\xi| - \xi \\ \xi \end{pmatrix} \tag{3.53}$$

and

$$K(S_t, \xi) = 0 \quad \forall S_t \in \mathbb{R}^2_+, \quad \xi \in \Xi,$$

where

$$S_t = \begin{pmatrix} X_t \\ Y_t \end{pmatrix}$$

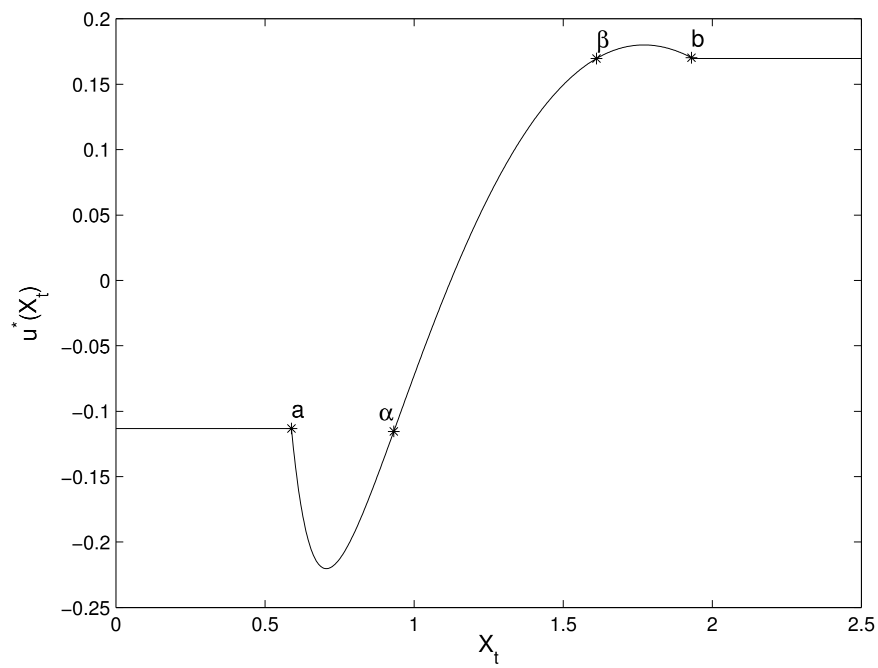

Figure 3.6: Optimal Interest Rate Control

such that

$$S_\tau = S_\tau^- + P(\xi_\tau).$$

The utility U_t of the investor is described by,

$$U(c_t) = \frac{c_t^\gamma}{\gamma},$$

where c_t is the consumption rate. The investor's consumption is withdrawn from the account of the risk-less asset, so that the dynamics of X_t may be defined as

$$dX_t = (rX_t - c_t)dt.$$

The choice of the consumption rate represents a continuous control for the investor. The combined control policy (c, ν) is said to be admissible if $0 \leq c \leq c_u$[8] and the state remains

[8] This constraint was imposed by Chancelier et al. [2002] for numerical reasons.

in the solvency region, denoted by $(c, \nu) \in \mathcal{W}$. The objective of the investor is to select the optimal policy for the both the impulse and stochastic control such that they maximize the expected discounted utility,

$$V(x,y) = \sup_{(c,\nu)\in\mathcal{W}} E\left[\int_0^\infty e^{-\delta t} U(c_t) dt \middle| X_0 = x, Y_0 = y\right],$$

where δ represents the agent's discount rate. As discussed in Section 3.2 the value function representing the maximized expected discounted flow of benefits will satisfy the quasi-variational inequalities in (3.44)-(3.45) where the optimal stochastic control is determined by (3.43). In this case the optimal consumption choice will be

$$c^* = \left[\frac{\partial V(x,y)}{\partial x}\right]^{\frac{1}{\gamma-1}},$$

within the interior $0 \leq c^* \leq c_u$.

As noted previously, the intervention operator $\mathcal{M}$ in (3.44) presents a challenge in terms of efficient implementation. Given the definition of the control's impact on the state variables in (3.53), the intervention operator may be interpreted as

$$\mathcal{M}V(X,Y) = \max_{\xi} V(X - k - \lambda|\xi| - \xi, Y + \xi),$$

subject to the solvency constraint that X and Y remain greater than or equal to zero. Therefore the intervention operator is a simple line search for argmax ξ. This procedure still represents a computationally costly process as the optimization problem needs to be solved for every nodal point in the approximation each time the EVNCP is evaluated. As discussed in Section 3.2 we reduce the complexity of the problem by discretizing the intervention operator as

$$\mathcal{M}V(X,Y) \approx \max\left(\Phi^M(X,Y)c\right), \tag{3.54}$$

where c is the vector of approximating coefficients and $\Phi^M(X,Y)$ is the family of approximating functions evaluated at M uniformly space points on the line that defines the admissible range for ξ. This basis matrix may be be pre-computed so that the intervention operator is reduced to a matrix multiplication operation and a simple maximum.

In order to solve the problem Chancelier et al. [2002] represent the combined stochastic and impulse control as the limit of an iterative series of combined stochastic control and optimal stopping problems. To solve the combined stochastic control and optimal stopping problem a policy iteration approach is combined with a finite difference approximation scheme for the value function. This method provided a solution that was computationally burdensome and numerically unstable. The technique presented in (3.2) is able to solve this problem in an efficient and accurate manner.

In order to implement their solution Chancelier et al. [2002] assume zero Neumann boundary conditions such that

$$\frac{\partial V(L,Y)}{\partial X} = 0 \tag{3.55}$$

and

$$\frac{\partial V(X,L)}{\partial Y} = 0. \tag{3.56}$$

where L represents the artificial boundary of the state space. As noted by Chancelier et al. [2002] this limits the accuracy of their result to a subset of the state space $[0, L] \times [0, L]$. In an attempt to limit the influence of imposing such a condition at an artificial boundary, we consider a transformation of variables mapping the true state space from $[0, \infty) \times [0, \infty)$ to $[0, 1] \times [0, 1]$. Letting $v(x, y) = V(X, Y)$ where

$$x = \frac{X}{a + X},$$

and

$$y = \frac{Y}{a + Y},$$

where a is a constant. Given this transformation the quasi-variational inequality defining the solution may be rewritten as

$$v(x, y) \geq \mathcal{M}v(x, y) \tag{3.57}$$

Table 3.4: Parameters for the Portfolio Management Example

Parameter	Description	Value
α	Risky asset drift rate	0.11
σ	Risky asset diffusion rate	0.3
r	Risk-less asset interest rate	0.07
δ	Discount rate	0.1
γ	Utility function parameter	0.3
k	Fixed transaction cost	0.05
λ	Proportional transaction cost	0.1
c_u	Maximum consumption rate	100

and

$$\delta v(x,y) \geq \frac{c^\gamma}{\gamma} + \left[rx(1-x) - \frac{c(1-x)^2}{a}\right]\frac{\partial v}{\partial x} + (\mu - \sigma^2 y)y(1-y)\frac{\partial v}{\partial y} + \frac{1}{2}\sigma^2 y^2(1-y)^2\frac{\partial^2 v}{\partial y^2}. \tag{3.58}$$

The parameters used in this example are the same as Chancelier et al. [2002] and are presented in Table 3.4. The family of approximating functions used are piecewise linear with 200 breakpoints in both directions of the x and y state space. We use centered finite difference methods in the interior of the state space and single sided finite difference methods on the boundaries. The transformation parameter a is set to 100. In the Newton algorithm the convergence criteria for the $\mathcal{L}^2$ norm of the residual is set to 1^{-8}. In evaluating the approximation of the intervention operator as described in (3.54) a value of $M = 300$ is used for each (x, y) node.

The optimal policy for the impulse control is presented in Figure 3.7. When in the “no transaction” region the investor should allow the state of the system to evolve uncontrolled until it hits one of the trigger boundaries. At this point it is optimal for the investor to exert an impulse that moves the state to the optimal target. In the case were the investor is initially outside the “no transaction” region at $t = 0$ it is optimal for to immediately move the state to within the region. As may be noted the transformation of variables allows the solution to be valid over the entire space $[0, 100] \times [0, 100]$ unlike the results presented in Chancelier et al. [2002], which are suggested to be valid only over a subset such as $[0, 50] \times [0, 50]$.

The value function is presented in Figure 3.8 and the optimal consumption rate is

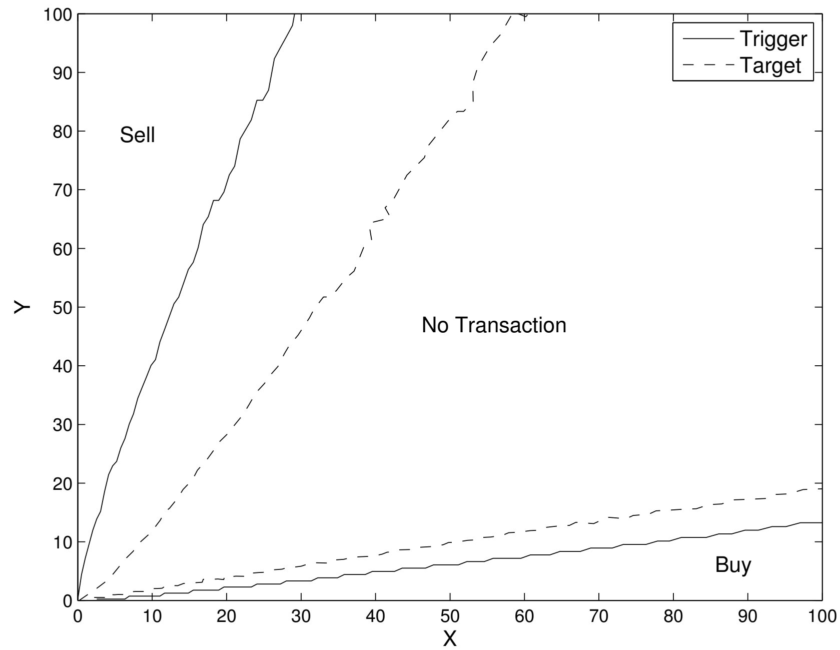

Figure 3.7: Impulse Control Boundaries for Portfolio Management Example

presented in Figure 3.9. We note that value of consumption is equal to zero when $X = 0$, due to the fact that strictly positive consumption requires the value of X to decrease while the solvency constraint prevents X from being negative. However, if the stock of the risky asset is greater than zero than it will be optimal for the agent to sell a portion of stock returning them to a part of the state space where consumption is greater than zero. It also worth noting that this situation would only occur if the agent were to start at this point in the state space, as the optimal policy for the impulse control would ensure that portfolio does not reach this point.

3.4 Concluding Remarks

With the growing interest in complex stochastic optimal control models that better accommodate reality, comes a need for efficient and accurate numerical solutions to such problems. This paper presents such a framework for the class of problems that combine opti-

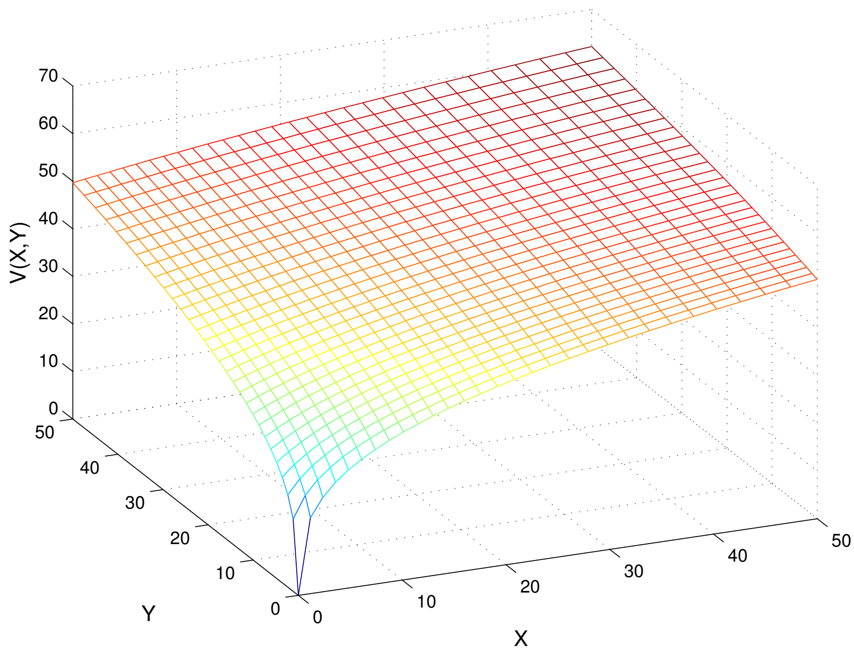

Figure 3.8: Portfolio Value

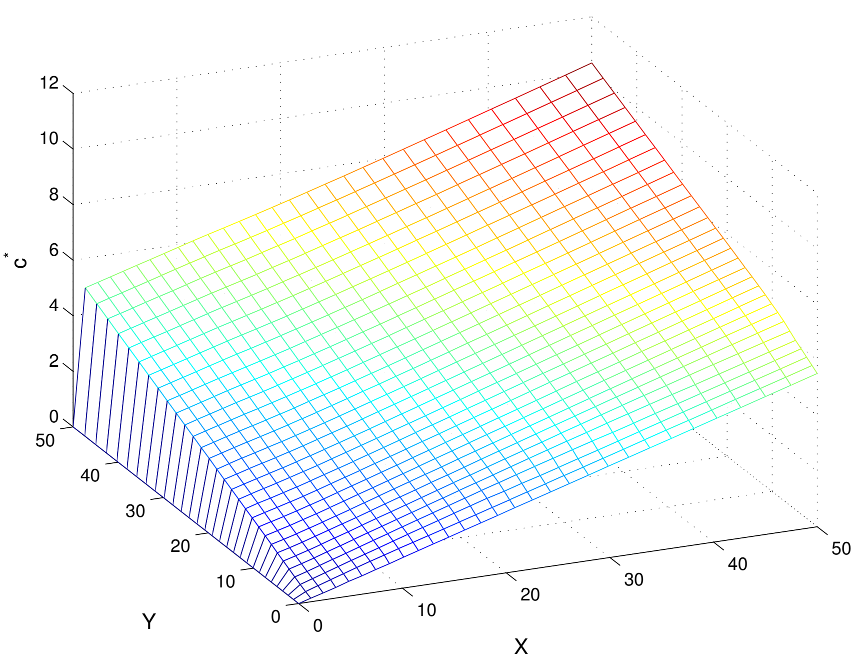

Figure 3.9: Optimal Consumption Rate

mal switching models with additional stochastic controls. The techniques discussed provide an attractive method of solving small to medium scale problems. As is the issue with multi-dimensional projection methods in general, the procedure suffers from the common curse of dimensionality. However, the techniques we present may be combined with other techniques for reducing the complexity of the quasi-variational inequalities such as transformation for variables and explicit finite differences methods for deterministic state variables.[9] Despite this shortcoming, the methods discussed are capable of handling the models currently being presented within the literature and can deal with additional complexity as is the case with the application in Chapter 2.

Based on the examples considered we find that there is a significant error associated with discretizing continuous controls in order to reduce the problem to that of only one class of control variable. To represent accurate approximations to the continuous control such techniques would require a level of discretization that produces a problem which is much more computationally burdensome than that with the continuous control. In fact, in the included resource extraction example the time required to compute the solution for the case with the continuous control was only slightly greater than that of a two regime switching model where the continuous control was held at a fixed value. We also find that combined impulse and stochastic control problems may easily be handled by either the direct EVNCP representation of the quasi-variational inequalities or through converting the problem to a combined optimal switching and stochastic control problem.

[9]See Chapter 2 for an example of handling a deterministic stock variable and Brekke and Oksendal [1994] for an example of variable transformations that reduces a problem's dimensionality.

Chapter 4

Efficient Estimation of One Factor Diffusion Models with Multivariate Gauss Hermite Quadrature

Continuous time Ito processes provide a convenient approach to modeling the dynamics of stochastic state variables and have therefore become increasingly popular in finance and economics. It is common to describe such a process by its stochastic differential equation (SDE) of the form

$$dX_t = a(X_t, \theta)dt + b(X_t, \theta)dW_t, \tag{4.1}$$

where $X(t_0) = X_0$, W is a standard Brownian motion, and $\theta \in \Theta \subset \mathbb{R}^p$. This paper focuses on univariate time homogenous processes where the drift and diffusion functions, $a(X_t, \theta)$ and $b(X_t, \theta)$, are known except for a vector of parameters θ.[1] The estimation of these unknown parameters, θ, proves to be a nontrivial problem for many applications due the discrete nature of data observations and a lack of known conditional transition densities. Given the wide use of SDEs a lot of effort has been put in to developing methods that provide an efficient estimate for θ using discrete observations. With much of this research

[1] Though the explicit focus of this paper is on univariate time homogenous processes the approach developed may be expanded to the multivariate and time heterogenous cases with some work.

there exists a trade off between computational accuracy and efficiency, in addition to ease of implementation. This paper presents a new estimation technique that provides accuracy with a computationally efficient algorithm that is relatively simply to apply.

The ideal estimator for the unknown parameter vector θ, is of course, the maximum likelihood estimate (MLE), such that

$$\hat{\theta} = \arg\max \sum_{t=1}^{T-1} \ln\left[f(X_{t+\Delta}; X_t, \theta)\right],$$

where Δ is the time step and $f(X_{t+\Delta}; X_t, \theta)$ is the conditional transition density for the process. However, the explicit conditional transition density is only known for a small number of specific processes. As a result much of the research on estimating θ has been focused on deriving efficient approximations for the conditional transition density.

The most common approach is to assume that the conditional transition density is Gaussian such that

$$X_{t+\Delta}|X_t \sim N\left[\mu(X_t, \theta, \Delta), \sigma^2(X_t, \theta, \Delta)\right].$$

where the moments are computed using a discrete approximation to (4.1). The simplest form of this discretization is the first order (Euler) approximation where

$$X_{t+\Delta} = X_t + a(X_t, \theta)\Delta + b(X_t, \theta)\sqrt{\Delta}\epsilon, \tag{4.2}$$

and $\epsilon \sim N(0, 1)$. Florens-Zmirou [1989] has shown that under mild conditions the quasi maximum likelihood estimate (QMLE), obtained with the first order Euler approximation, will converge to the true MLE as the sampling interval goes to zero. However for common frequencies found in economics, such as weekly and monthly, experiments have shown that such estimates will exhibit significant bias (see e.g. Ait-Sahalia [1999], Durham and Gallant [2002]). Alternative approaches to compute the moments of the Gaussian transition density have fared better in tests. Shoji and Ozaki [1998] present a linearization of the SDE resulting in an Ornstein-Uhlenbeck process that has a Gaussian transition density for which the true moments are known. Kessler [1997] uses a higher order Ito-Taylor expansion of the SDE in order to approximate the moments for the Gaussian conditional transition density.

The trouble with methods based on a discrete approximation of (4.1) is that they all require the sampling interval to go to zero in order to obtain convergence to the true conditional transition density. As an alternative the forward Kolmogorov (Fokker-Planck) partial differential equation, which is satisfied by the transition density, may be numerically solved for an approximation of the true density (see e.g. Lo [1988]). The method of Ait-Sahalia [2002] and Bakshi and Ju [2005] introduces a closed form approximation of the conditional transition density based on an Ito-Taylor Hermite expansion around a Gaussian density. Though this technique still requires the sampling interval to approach zero for finite expansions the approximation may made arbitrarily close to the true transition density through the inclusion of additional correction terms.

Though attractive due to a lack of reliance on small sampling intervals, such methods are uncommon amongst practitioners due to either poor numerical accuracy or difficulty of implementation. As a result, research on the use of discrete approximations has persisted, primarily in the form of the simulated maximum likelihood estimate (SMLE). First introduced by Pedersen [1995] the technique incorporates points between each pair of observations and then integrates out the unknown states (see also Brandt and Santa-Clara [2002]). Noting that the first order discrete approximation in (4.2) will only be accurate for small time steps the SMLE method incorporates sub-intervals between the points t and $t + \Delta$, where $t = \tau_0 < \tau_1 < ... < \tau_I = t + \Delta$. The number of sub-intervals, I, is chosen so that the Gaussian assumption for the conditional transition density will be accurate for the smaller time step Δ/I. However, the points $X(\tau_1), ..., X(\tau_{I-1})$ will be unobserved. Fortunately the Markovian nature of the process allows these observations to be integrated out through an application of the Chapman-Kolmogorov equation, such that

$$f^{(I)}(X_t; X_{t+\Delta}, \theta) = \int \prod_{i=0}^{I-1} f(z_{i+1}; z_i, \theta) d\lambda(z_1, ..., z_{I-1}), \tag{4.3}$$

where λ is the Lebesgue measure, $z_0 = X_t$ and $z_I = X_{t+\Delta}$. As the name suggests, the SMLE approach performs the integration using Monte Carlo simulations. Pedersen [1995] shows that under mild conditions the approximation $f(X_t; X_{t+\Delta}, \theta) \approx f^{(I)}(X_t; X_{t+\Delta}, \theta)$ will converge as I becomes large.

The original version of the SMLE method was attractive because the technique is easily applied to any process with an explicit SDE. However, the method is computationally costly as it requires a large number of sample paths generated over a large number of sub-intervals for each pair of observations. Furthermore even with a high value for I and a large number of sample paths the method is not as computationally accurate as alternative approaches. Durham and Gallant [2002] provided a substantial update to the SMLE technique by utilizing moments derived from higher order approximations such as those introduced by Shoji and Ozaki [1998] and Kessler [1997], in order to improve the approximation of the sub-densities. In addition they incorporated techniques from the Markov Chain Monte Carlo (MCMC) methods of Eraker [2001], Jones [1999], and Elerian et al. [2001], to improve the numerical integration of (4.3).[2] As a result of these improvements they are able to increase the accuracy of the SMLE. However, the SMLE approach is still computationally burdensome due to the choice of Monte Carlo integration and the relatively large number of sub-intervals required to obtain an accurate approximation of the conditional transition density.

We purpose an alternative method based on the approximation of (4.3), in which the unobserved states are integrated out using multivariate Gauss Hermite quadrature as opposed to Monte Carlo simulation. The benefit of numerical quadrature over Monte Carlo integration is that for low to moderate values of I the integral may be computed to an arbitrary level of accuracy in a significantly shorter period of time. This provides the opportunity to compute approximations of the conditional transition density for multiple time steps which may then be combined using Richardson extrapolation in order to further improve the approximation. Given the accuracy of the numerical integration and the benefits of Richardson extrapolation we are able to achieve the same computational accuracy as the updated SMLE using only a one dimensional integral as opposed to the fifteen dimensional integral required in the SMLE. The method therefore provides a significant improvement in computational efficiency without any loss of accuracy.

[2] We refer the reader to Hurn et al. [2007] for an overview of such likelihood based estimation techniques, in addition to non-likelihood based estimation methods for SDEs.

The approximation for the conditional transition density presented in (4.3) is associated with three sources of error in its computation. The first two stem from the approximation of the density and its conditional moments when computing the sub-density. Whereas the third source of error arises in the numerical integration used in computing the Chapman-Kolmogorov equation. The paper proceeds by first discussing methods of reducing the error associated with computing the sub-density in Section 4.1, followed by the presentation in Section 4.2 of an accurate method for integrating out the sub-intervals. In Section 4.3 we examine a number of experiments to test the accuracy of the new estimator and in Section 4.4 we provide concluding remarks.

4.1 Approximating the Sub-Density

The approximation for the conditional transition density presented in (4.3) still requires a transition density to be defined for the subintervals. For this purpose it is possible to utilize the first order Euler approximation in (4.2) such that the density for an observation at time $t + \Delta$, given the information at time t, is assumed to be Gaussian, where

$$f(X_{t+\Delta}; X_t, \theta) = \frac{1}{\sigma(X_t, \theta, \Delta)\sqrt{2\pi}} exp \left\{ -\frac{[X_{t+\Delta} - \mu(X_t, \theta, \Delta)]^2}{2\sigma^2(X_t, \theta, \Delta)} \right\},$$

where the mean, $\mu(X_t, \theta, \Delta)$, and variance, $\sigma^2(X_t, \theta, \Delta)$, are defined as

$$\mu(X_t, \theta, \Delta) = X_t + a(X_t, \theta)\Delta,$$

and

$$\sigma^2(X_t, \theta, \Delta) = b^2(X_t, \theta)\Delta.$$

As shown by Florens-Zmirou [1989] this approximation will converge to the true conditional transition density as Δ approaches zero. It is worth noting that the error of this approximation will be of order Δ (see Bally and Talay [1995]). Therefore in theory this approximation for the sub-densities will be accurate if the number of sub-intervals, I, is large enough so that the step size is made sufficiently small. However in practice the step size required for the Euler method to provide an accurate approximation of the conditional transition density would requires a relatively high number of sub-intervals to be included. As a result (4.3) includes a high dimensional integral that must be solved for each pair of observations. Therefore there exists a benefit to considering more accurate approximations for the conditional transition density, such that an integral of a lower dimension may be used.

This section presents a set of three methods for increasing the accuracy of the sub-density approximation while maintaining the assumption that the conditional transition density will be Gaussian. The first technique involves transforming the data so that it appears more Gaussian. The second two methods provide improvements in the approximation of the conditional transition density by utilizing second order approximations of the SDE including a second order Milstein approach and the local linearization method of Shoji and Ozaki [1998].

4.1.1 Euler with Constant Diffusion

There is evidence to suggest that transforming the SDE to a process with a constant diffusion term will reduce the error associated with the approximation (see Bally and Talay [1995]), as the conditional transition density for a constant diffusion process is more Gaussian in nature. To obtain a unit diffusion process consider the transformation

$$Y_t \equiv \gamma(X_t, \theta) = \int^{X_t} b^{-1}(u, \theta) du. \tag{4.4}$$

The process governing the dynamics of Y may be obtained via Ito's lemma and will have the general form

$$dY_t = \alpha(Y_t, \theta)dt + \beta dW_t, \tag{4.5}$$

where $\beta = 1$ for the transformation in (4.4).[3] Under this transformation the Euler first order approximation for the first two moments of the Gaussian conditional transition density,

[3] We include β as a reminder that it is not required that the diffusion be equal to one as is the case with the transformation in (4.4). In fact from a computational standpoint transformations that do not explicitly depend on the parameters θ yet produce a constant diffusion process will be preferred.

$f(Y_{t+\Delta}; Y_t, \theta)$ are

$$\mu(Y_t, \theta, \Delta) = Y_t + \alpha(Y_t, \theta)\Delta,$$

and

$$\sigma^2(Y_t, \theta, \Delta) = \beta^2 \Delta.$$

The conditional transition density approximation using the Euler method for the constant diffusion SDE will provide an improvement in the accuracy of the approximation, though it will remain of order Δ.

4.1.2 Second Order Milstein Approximation

To examine the possibility of higher order approximations for the sub-density we turn to the second order Milstein approximation for the non-constant diffusion SDE in (4.1).[4] The SDE may be rewritten in integral form as

$$X_{t+\Delta} = X_t + \int_0^{\Delta} a(X_{t+\tau}, \theta) d\tau + \int_0^{\Delta} b(X_{t+\tau}, \theta) dW_{t+\tau}. \tag{4.6}$$

The Milstein approximation utilizes the fact that (4.6) may be expressed as

$$X_{t+\Delta} = X_t + \int_0^{\Delta} \left[a(X_t, \theta) + \int_0^{\tau} da(X_{t+s}, \theta) \right] d\tau + \int_0^{\Delta} \left[b(X_t, \theta) + \int_0^{\tau} db(X_{t+s}, \theta) \right] dW_{t+\tau}. \tag{4.7}$$

Applying Ito's lemma to the drift and diffusion functions of (4.1) results in

$$da_t(X_t, \theta) = \left[a'(X_t, \theta) a(X_t, \theta) + \frac{b^2(X_t, \theta) a''(X_t, \theta)}{2} \right] dt + a'(X_t, \theta) b(X_t, \theta) dW_t, \tag{4.8}$$

and

$$db_t(X_t, \theta) = \left[b'(X_t, \theta) a(X_t, \theta) + \frac{b^2(X_t, \theta) b''(X_t, \theta)}{2} \right] dt + b'(X_t, \theta) b(X_t, \theta) dW_t.$$

[4]For a detailed discussion of higher oder SDE approximations we refer the interested reader to Kloeden and Platen [1992].

Given these dynamics for the drift and diffusion functions and holding X constant at X_t will yield

$$\int_0^\Delta da(X_{t+\tau},\theta) \approx \left[a'(X_t,\theta)a(X_t,\theta)+\frac{1}{2}b^2(X_t,\theta)a''(X_t,\theta)\right]\Delta + a'(X_t,\theta)b(X_t,\theta)(W_{t+\tau}-W_t), \tag{4.9}$$

and

$$\int_0^\Delta db(X_{t+\tau},\theta) \approx \left[b'(X_t,\theta)a(X_t,\theta)+\frac{1}{2}b^2(X_t,\theta)b''(X_t,\theta)\right]\Delta + b'(X_t,\theta)b(X_t,\theta)(W_{t+\tau}-W_t). \tag{4.10}$$

Substituting (4.9) and (4.10) into (4.7) yields

$$\begin{aligned} X_{t+\Delta} \approx\ & X_t + a(X_t,\theta)\Delta + \tfrac{1}{2}\left[a'(X_t,\theta)a(X_t,\theta)+\tfrac{1}{2}b^2(X_t,\theta)a''(X_t,\theta)\right]\Delta^2 \\ & +a'(X_t,\theta)b(X_t,\theta)\textstyle\int_0^\Delta (W_{t+\tau}-W_t)d\tau + b(X_t,\theta)\int_0^\Delta dW_{t+\tau} \\ & +\left[b'(X_t,\theta)a(X_t,\theta)+\tfrac{1}{2}b^2(X_t,\theta)b''(X_t,\theta)\right]\textstyle\int_0^\Delta \tau dW_{t+\tau} \\ & +b'(X_t,\theta)b(X_t,\theta)\textstyle\int_0^\Delta (W_{t+\tau}-W_t)dW_{t+\tau}. \end{aligned} \tag{4.11}$$

Noting that

$$\int_0^\Delta dW_{t+\tau} \approx (W_{t+\Delta}-W_t),$$

$$\int_0^\Delta (W_{t+\tau}-W_t)dW_{t+\tau} = \frac{1}{2}\left[(W_{t+\Delta}-W_t)^2-\Delta\right],$$

and

$$\int_0^\Delta (W_{t+\tau}-W_t)d\tau = \int_0^\Delta \tau dW_{t+\tau} = Z_{t+\Delta}-Z_t,$$

where $Z_{t+\Delta} \sim N(Z_t, \Delta^3/3)$, (4.11) may be rewritten as

$$\begin{aligned} X_{t+\Delta} \approx\ & X_t + a(X_t,\theta)\Delta + \tfrac{1}{2}\left[a'(X_t,\theta)a(X_t,\theta)+\tfrac{1}{2}b^2(X_t,\theta)a''(X_t,\theta)\right]\Delta^2 \\ & +a'(X_t,\theta)b(X_t,\theta)(Z_{t+\Delta}-Z_t) + \tfrac{1}{2}b'(X_t,\theta)b(X_t,\theta)\left[(W_{t+\Delta}-W_t)^2-\Delta\right] \\ & +\left[b'(X_t,\theta)a(X_t,\theta)+\tfrac{1}{2}b^2(X_t,\theta)b''(X_t,\theta)\right]\left[\Delta(W_{t+\Delta}-W_t)-(Z_{t+\Delta}-Z_t)\right] \\ & +b(X_t,\theta)(W_{t+\Delta}-W_t). \end{aligned}$$

Given that the $\mathrm{Cov}(W_{t+\Delta}, Z_{t+\Delta}|W_t, Z_t) = \Delta^2/2$, we can replace the random variable $Z_{t+\Delta} - Z_t$ with $\Delta/2(W_{t+\Delta} - W_t) + \tilde{W}_t$ where $\tilde{W}_t \sim N(0, \Delta^3/12)$, such that

$$\begin{aligned} X_{t+\Delta} \approx\ & X_t + a(X_t,\theta)\Delta + \tfrac{1}{2}\left[a'(X_t,\theta)a(X_t,\theta) + \tfrac{1}{2}b^2(X_t,\theta)a''(X_t,\theta)\right]\Delta^2 \\ & + \tfrac{\Delta}{2}\left[a'(X_t,\theta)b(X_t,\theta) + b'(X_t,\theta)a(X_t,\theta) + \tfrac{1}{2}b^2(X_t,\theta)b''(X_t,\theta)\right](W_{t+\Delta} - W_t) \\ & + \left[a'(X_t,\theta)b(X_t,\theta) - b'(X_t,\theta)a(X_t,\theta) - \tfrac{1}{2}b^2(X_t,\theta)b''(X_t,\theta)\right]\tilde{W}_t \\ & + \tfrac{1}{2}b'(X_t,\theta)b(X_t,\theta)\left[(W_{t+\Delta} - W_t)^2 - \Delta\right] + b(X_t,\theta)(W_{t+\Delta} - W_t). \end{aligned} \tag{4.12}$$

The presence of the term $(W_{t+\Delta} - W_t)^2$ suggests that the second order approximation of the conditional transition density for a non-constant diffusion process is not normal. In fact it may be shown that the correct density is a convolution of a normal density and a non-central chi-square density.

If we instead consider the constant diffusion transformed process in (4.5), (4.12) will reduce to

$$\begin{aligned} Y_{t+\Delta} \approx\ & Y_t + \alpha(Y_t,\theta)\Delta + \tfrac{1}{2}\left[\alpha'(Y_t,\theta)\alpha(Y_t,\theta) + \tfrac{1}{2}\beta^2\alpha''(Y_t,\theta)\right]\Delta^2 \\ & + \beta\left[\left(1 + \tfrac{\alpha'(Y_t,\theta)\Delta}{2}\right)(W_{t+\Delta} - W_t) + \alpha'(Y_t,\theta)\tilde{W}_t\right]. \end{aligned}$$

Therefore in the constant diffusion case the second order approximation for the conditional transition density will be normal with conditional mean and variance

$$\mu(Y_t,\theta,\Delta) = Y_t + \alpha(Y_t,\theta)\Delta + \frac{1}{2}\left[\alpha'(Y_t,\theta)\alpha(Y_t,\theta) + \frac{1}{2}\beta^2\alpha''(Y_t,\theta)\right]\Delta^2,$$

and

$$\sigma^2(Y_t,\theta,\Delta) = \beta^2\left[1 + \alpha'(Y_t,\theta)\Delta + \frac{(\alpha'(Y_t,\theta))^2}{3}\Delta^2\right]\Delta.$$

4.1.3 Shoji and Ozaki

Shoji and Ozaki [1998] purpose an alternative second order approximation for the conditional transition density that is based on a local linearization. Henceforth we denote this approach as the Shoji & Ozaki method. Given the constant volatility process in (4.5), applying Ito's lemma to the drift function results in

$$d\alpha_t(Y_t,\theta) = \left[\alpha'(Y_t,\theta)\alpha(Y_t,\theta) + \frac{\beta^2\alpha''(Y_t,\theta)}{2}\right] dt + \alpha'(Y_t,\theta)\beta dW_t. \tag{4.13}$$

By assuming that for a small time step, Δ, the first two derivatives of the drift are constant, the process in (4.13) is an Ornstein-Uhlenbeck process, whose conditional transition density is known to be Gaussian. It may be seen that the first moment of the process solves the ordinary differential equation (ODE)

$$\frac{dE[\alpha(Y_{t+\Delta},\theta)|Y_t]}{d\Delta} = \alpha'(Y_t,\theta)E[\alpha(Y_{t+\Delta},\theta)|Y_t] + \frac{\beta^2\alpha''(Y_t,\theta)}{2},$$

subject to the boundary condition $E[\alpha(Y_t,\theta)|Y_t] = \alpha(Y_t,\theta)$. The solution to the ODE is

$$E[\alpha(Y_{t+\Delta},\theta)|Y_t] = \left[e^{\alpha'(Y_t,\theta)\Delta} - 1\right] \frac{\beta^2\alpha''(Y_t,\theta)}{2\alpha'(Y_t,\theta)} + e^{\alpha'(Y_t,\theta)\Delta}\alpha(Y_t,\theta).$$

Given this expectation for the local linearization of the drift function the expectation of the process, Y, will be:

$$\begin{aligned} E[Y_{t+\Delta}|Y_t,\theta] = & \ Y_t + \textstyle\int_0^\Delta E[\alpha(Y_{t+\tau},\theta)|Y_t]d\tau \\ = & \ Y_t + \tfrac{e^{\alpha'(Y_t,\theta)\Delta}-1}{\alpha'(Y_t,\theta)}\alpha(Y_t,\theta) + \left[\tfrac{e^{\alpha'(Y_t,\theta)\Delta}-1}{\alpha'(Y_t,\theta)} - \Delta\right] \tfrac{\beta^2\alpha''(Y_t,\theta)}{2\alpha'(Y_t,\theta)}. \end{aligned}$$

From the Kolmogorov backward equation it may be seen that the conditional variance of the process will be

$$\begin{aligned} Var[Y_{t+\Delta}|Y_t,\theta] = & \ \left[\textstyle\int_0^\Delta e^{2\alpha'(Y_t,\theta)\tau}d\tau\right]\beta^2 \\ = & \ \tfrac{e^{2\alpha'(Y_t,\theta)\Delta}-1}{2\alpha'(Y_t,\theta)}\beta^2. \end{aligned}$$

Therefore, the mean and variance of the Gaussian conditional transition density for $Y_{t+\Delta}$ given Y_t are $\mu(Y_t,\theta,\Delta) = E[Y_{t+\Delta}|Y_t,\theta]$ and $\sigma^2(Y_t,\theta,\Delta) = Var[Y_{t+\Delta}|Y_t,\theta]$ respectively.

4.2 Integration Over Sub-Intervals

Given a representation for the sub-density the approximation for the conditional transition density $f(X_t; X_{t+\Delta}, \theta) \approx f^{(I)}(X_t; X_{t+\Delta}, \theta)$ is defined as in (4.3). Due to the assumption of normality for the sub-densities this integration may be approximated using Gauss Hermite quadrature such that

$$f^{(I)}(X_{t+\Delta}; X_t, \theta) \approx \sum_{n=1}^{N} \omega_n \phi(X_{t+\Delta}; z_{I-1,n}, \theta),$$

where $\omega_1, ..., \omega_N$ are the set of quadrature weights, ϕ is the normal density function, and

$$z_{i,n} = \mu\left(z_{i-1,n}, \theta, \frac{\Delta}{I}\right) + \sigma\left(z_{i-1,n}, \theta, \frac{\Delta}{I}\right) u_{i,n},$$

where $u_{i,1}, ..., u_{i,N}$ are the set of quadrature nodes for subinterval i and $z_0 = X_t$. Based on our testing we suggest the use of standard Gauss Hermite quadrature for the one dimensional integral in the $f^{(2)}(X_{t+\Delta}; X_t, \theta)$ approximation and use of the tensor product of one dimensional rules for the multivariate cases of $I \geq 3$ (see Miranda and Fackler [2002]). For this application sparse Gauss Hermite quadrature methods appear to have poor performance requiring the use of a greater number of nodes, than the tensor product, in order to achieve a given level of accuracy.

The use of numerical quadrature to perform the integration as opposed to Monte Carlo simulations, as is the case with the SMLE, allows for a significant improvement in both accuracy of the integral approximation and the computational speed for low to moderate values of I. As a result it is computationally cheap to evaluate $f^{(I)}(X_t; X_{t+\Delta}, \theta)$ for multiple values of I allowing for the implementation of Richardson extrapolation to further reduce the order of the error associated with the conditional transition density approximation. For example, consider an approximation of the transition density with a step size of Δ and error of order Δ^n, such that

$$f^{(1)}(X_{t+\Delta}; X_t, \theta) = f(X_{t+\Delta}; X_t, \theta) + a\Delta^n + O(\Delta^{n+1}).$$

An approximation of the transition density with a step size of Δ/I may be similarly represented as

$$f^{(I)}(X_{t+\Delta};X_t,\theta) = f(X_{t+\Delta};X_t,\theta) + a\left(\frac{\Delta}{I}\right)^n + O\left(\Delta^{n+1}\right).$$

It is possible to determine coefficients, λ_1 and λ_2, such that the approximation of conditional transition density

$$f(X_{t+\Delta};X_t,\theta) \approx \lambda_1 f^{(1)}(X_{t+\Delta};X_t,\theta) + \lambda_2 f^{(I)}(X_{t+\Delta};X_t,\theta),$$

is of order Δ^{n+1}. It is easily seen that in this example the weights must satisfy

$$\lambda_1 + \lambda_2 = 1 \tag{4.14}$$

and

$$\lambda_1 + \frac{\lambda_2}{I^n} = 0. \tag{4.15}$$

Therefore solutions for the coefficients may be obtained by solving the system presented in (4.14)-(4.15), so that

$$\lambda_1 = \frac{1}{1-I^n}, \qquad \lambda_2 = \frac{I^n}{I^n-1}.$$

This method may be easily generalized to a framework for approximations of the conditional transition density $f^{(1)}$, $f^{(2)}$, ..., $f^{(I)}$ each of order Δ^n, $\left(\frac{\Delta}{2}\right)^n$, .., $\left(\frac{\Delta}{I}\right)^n$ respectively. The new approximation is given by

$$f(X_{t+\Delta};X_t,\theta) \approx \sum_{i=1}^{I} \lambda_i f^{(i)}(X_{t+\Delta};X_t,\theta), \tag{4.16}$$

where the weights satisfy

$$\sum_{i=1}^{I} \lambda_i = 1 \tag{4.17}$$

and the $I-1$ conditions

$$\sum_{i=1}^{I} \frac{\lambda_i}{(i)^{n+h}} = 0 \qquad h = 0, ..., I-2. \tag{4.18}$$

The square linear system in (4.17)-(4.18) may easily be solved for the extrapolation coefficients. As a result the new approximation of the conditional transition density in (4.16) will be of order Δ^{n+I-1}.

The Euler approximation is of weak order Δ, while the Milstein approximation along with the one developed by Shoji and Ozaki [1998] are of weak order Δ^2. As demonstrated in Section 4.3 the benefits associated with the application of Richardson extrapolation are substantial. It is worth noting that the ability of Richardson extrapolation to reduce the order of the approximation error is based on the assumption that the integration in $f^{(I)}(X_{t+\Delta}; X_t, \theta)$ is computed to a high degree of accuracy. If this is not the case the application of Richardson extrapolation has the possibility of reducing the accuracy of the conditional transition density approximation. To see this consider the example of $\lambda_1 f^{(1)}(X_{t+\Delta}; X_t, \theta) + \lambda_2 f^{(2)}(X_{t+\Delta}; X_t, \theta)$. Given that integration in $f^{(2)}(X_{t+\Delta}; X_t, \theta)$ is numerically approximated it will be the case that

$$\lambda_1 f^{(1)}(X_{t+\Delta}; X_t, \theta) + \lambda_2 f^{(2)}(X_{t+\Delta}; X_t, \theta) = f(X_{t+\Delta}; X_t, \theta) + \lambda_2 \epsilon + O(\Delta^{n+1}),$$

where ϵ represents the error associated with the numerical integration. In order for Richardson extrapolation to be beneficial the integration error must be small enough that it does not dominate the error associated with the sub-density approximation. This is why Richardson extrapolation is not an effective technique for reducing the approximation error in the SMLE. The number of simulated paths required to achieve the needed accuracy in the Monte Carlo integration would result in a significant computational burden. However, the utilization of Gauss Hermite quadrature in approximating (4.3) allows for relatively quick and accurate numerical integration making the implementation of Richardson extrapolation attractive.

With the use of Richardson extrapolation it is possible obtain negative conditional transition densities for particular sets of parameters and observations. Any application of the technique should check for this error and fall back to $f^{(\bar{I})}(X_{t+\Delta}; X_t, \theta)$ if encountered, where $\bar{I}$ represents the best approximation computed for use with Richardson extrapolation.

4.3 Numerical Examples

In order to examine the efficiency of the purposed conditional transition density approximation we consider a set of numerical examples. To gauge the method's numerical properties in addition to its performance in empirical analysis we study both simulated experiments and applications utilizing observed data. The first test of the new technique follows the setup of Ait-Sahalia [1999] in which the parameters of SDEs with known transition densities are estimated using monthly observations of the federal funds rate. This provides the ability to examine the efficiency of the new approximation, as compared to the MLE, in an application setting. We find that incorporating only a single sub-interval produces estimates that are only negligibly different from the true MLE. To further examine the properties of the new estimator we consider a set of benchmarks based on experiments using the Cox-Ingersol-Ross process. Again we find the new method to be both computationally accurate and efficient. We conclude this section with a a series of examples to showcase the performance of the approximation in a variety of complex cases.

4.3.1 Diffusion Models for Interest Rates

Ait-Sahalia [1999] presents a study of various models which have been utilized to describe the dynamics of short term interest rates. For our purpose we focus on a set of three models, for which explicit expressions for the true conditional transition densities are known. As such these models present an ideal basis to judge the performance of conditional transition densities in an empirical setting. The models of interest are the Ornstein-Uhlenbeck, Cox-Ingersol-Ross, and inverse square root process. In addition we examine the constant elasticity of variance model for which there is no closed form solution for the conditional transition density, but other accurate approximations are available for comparison of the estimates.

Ornstein-Uhlenbeck

The Ornstein-Uhlenbeck process is a mean reverting process whose SDE takes the form

$$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3 dW_t.$$

This process has a Gaussian transition density such that

$$f^{OU}(X_{t+\Delta}; X_t, \theta) = \sqrt{\frac{\theta_2}{\pi\gamma^2}} \exp\left\{ -\frac{\left[X_{t+\Delta} - \theta_1 - (X_t - \theta_1)e^{-\theta_2\Delta}\right]^2 \theta_2}{\gamma^2} \right\},$$

where $\gamma^2 = \theta_3^2\left(1 - e^{-2\theta_2\Delta}\right)$. It may be noted that the diffusion function is constant so no transformation is necessary.

Cox-Ingersol-Ross

The process introduced by Feller [1952] has been used in numerous financial applications, most notably is the application to short term interest rates by Cox et al. [1985]. As a result the process is commonly referred to as the Cox-Ingersol-Ross or CIR process. Its SDE has the form

$$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3\sqrt{X}dW. \tag{4.19}$$

The true conditional transition density of (4.19) is

$$f^{CIR}(X_{t+\Delta}; X_t, \theta) = ce^{-(u+v)} \left(\frac{v}{u}\right)^{q/2} I_q(2\sqrt{uv}), \tag{4.20}$$

where

$$c = \frac{2\theta_2}{\theta_3^2(1 - e^{-\theta_2\Delta})},$$

$$u = cX_t e^{-\theta_2\Delta},$$

$$v = cX_{t+\Delta},$$

$$q = \frac{2\theta_2\theta_1}{\theta_3^2} - 1,$$

and $I_q(\cdot)$ is a modified Bessel function of the first kind of order q. It may be seen that the transformed random variable, $Y = 2cX$, has a conditional transition density that is non-central chi-square with $4\theta_2\theta_1/\theta_3^2$ degrees of freedom and a non-centrality parameter $Y_t e^{-\theta_2\Delta}$.

In the cases where the process must be transformed to have a constant diffusion the transformation $Y_t = 2\sqrt{X_t}$ is used, such that

$$dY_t = \left[\left(2\theta_1\theta_2 - \frac{\theta_3^2}{2}\right)\frac{1}{Y_t} - \frac{\theta_2}{2}Y_t\right]dt + \theta_3 dW. \tag{4.21}$$

Inverse of Feller's Square Root Model

The inverse of Feller's square root process (ISR) has a specification of one over the CIR model. From Ito's Lemma the SDE is defined as

$$dX_t = X_t[\theta_2 - (\theta_3^2 - \theta_2\theta_1)X_t]dt + \theta_3 X_t^{3/2} dW_t.$$

The true conditional transition density is given by

$$f^{ISR}(X_{t+\Delta}; X_t, \theta) = \frac{1}{X_t^2} f^{CIR}\left(\frac{1}{X_{t+\Delta}}; \frac{1}{X_t}, \theta\right),$$

where f^{CIR} is the conditional transition density for the CIR process as defined in (4.20). In the cases where the process must be transformed to have a constant diffusion the transformation $Y_t = 2/\sqrt{X_t}$ is used. The form of the transformed process will equivalent to that of (4.21).

Constant Elasticity of Variance

The constant elasticity of variance (CEV) model proposed by Chan et al. [1992] has the form

$$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3 X_t^{\theta_4} dW_t.$$

The CIR process is a special case of the CEV process in which $\theta_4 = 1/2$. However, in the case of the more general CEV process there does not exist a closed form solution for the conditional transition density. The constant diffusion transformation is also dependent upon the unknown parameters, such that $Y_t = X_t^{1-\theta_4}/(1-\theta_4)$ where transformed SDE has the form

$$dY_t = \left\{ \theta_2[\gamma(Y_t, \theta)]^{-\theta_4} [\theta_1 - \gamma(Y_t, \theta)] - \frac{1}{2}\theta_3^2\theta_4[\gamma(Y_t, \theta)]^{\theta_4 - 1} \right\} dt + \theta_3 dW_t,$$

where

$$\gamma(Y_t, \theta) = [Y_t (1-\theta_4)]^{\frac{1}{1-\theta_4}} .$$

Results

Following the work of Ait-Sahalia [1999] parameter estimates for the three models are computed using a monthly sampling of the Federal Funds rate from January 1963 to December 1998. The observations are presented in Figure 4.1. To represent the new method we utilize the constant diffusion transformations of the models along with the mean and variance approximations generated from the second order Milstein specification discussed in Section 4.1.2. Using Richardson extrapolation we combine the estimates from the case using the standard step size, $I = 1$, and one where we introduce an additional sub-interval, $I = 2$. As such the integral in (4.3) is only over one dimension and as such is relatively quick to compute. In order to obtain the nodes and weights for the Gauss Hermite quadrature we utilize the standard one dimensional rule with "order" $M = 25$.[5] We denote this new method by $f^{(1)} + f^{(2)}$ to conserve on notation. In the subsequent section we explore in depth the properties of various forms of the new method. Here we demonstrate that in empirical applications even a low dimension form of the new approximation may provide accurate estimates

The parameter estimates for the new method, $f^{(1)} + f^{(2)}$, along with the MLE are presented in Table 4.1. For comparison we also include the estimates from the basic first order Euler approximation of Section 4.1 in addition to the Hermite expansion approximation presented in Ait-Sahalia [1999]. As may be seen, the Euler approximation provides poor performance in matching the MLE. However, the new method, even with only two

[5]The term order is used to represent the degree of polynomial which may be integrated exactly with the quadrature technique. An order of M corresponds to a polynomial of degree $2M - 1$.

sub-intervals, is able to produce estimates relatively close to the MLE. Given that only one dimensional integration needs to be performed in order to obtain these estimates they are relatively fast to compute. In fact during these experiments, using the new method with $f^{(1)} + f^{(2)}$ and the Milstein approximation took at most .01 seconds longer to compute a log-likelihood as compared to the method of Ait-Sahalia [1999] with K=2. The difference in the total time required to compute the estimates using the two methods will of course be based on the number of function evaluations required by the optimization algorithm, but in practice differed by less than half a second in most cases. This additional computational burden is quite small compared to the additional preparation time required to derive and program the complex Hermite expansions required by the method of Ait-Sahalia [1999]. Therefore given the simplicity of the technique presented in this paper, the overall time required to obtain parameter estimates for an additional model will be significantly less than that associated with utilizing the Hermite expansion method, and without a loss of accuracy.

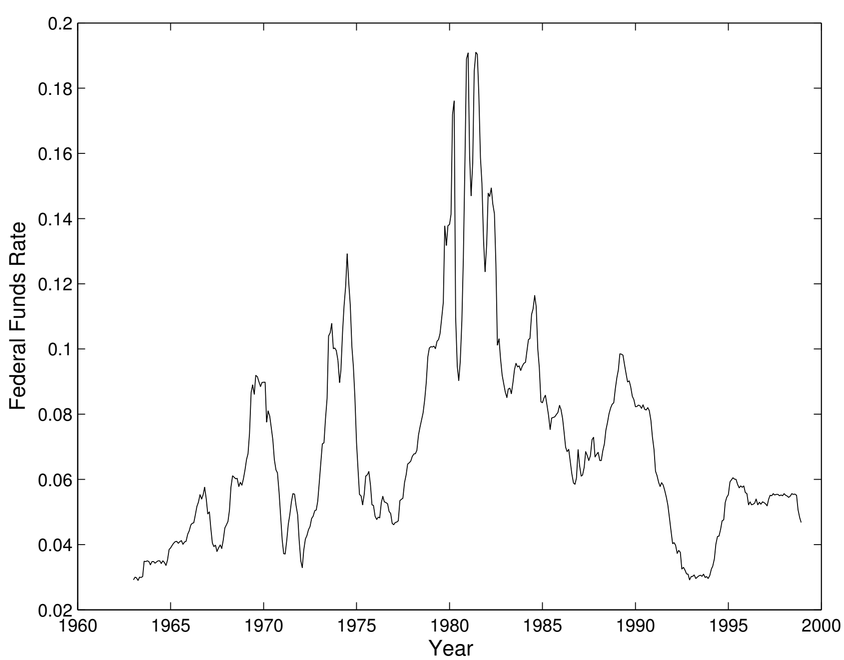

Figure 4.1: Monthly Federal Funds Rate, 1963-1998

Table 4.1: Maximum-Likelihood Estimates for the Monthly Federal Funds Rate

Model	**Euler**	**Ait Sahalia K=1**	**Ait Sahalia K=2**	$f^{(1)} + f^{(2)}$	**True Density**
$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3 dW_t$	$\theta_1 = 0.0717$ $\theta_2 = 0.2584$ $\theta_3 = 0.0222$	$\theta_1 = 0.0720$ $\theta_2 = 0.2575$ $\theta_3 = 0.0224$	$\theta_1 = 0.0717$ $\theta_2 = 0.2614$ $\theta_3 = 0.0224$	$\theta_1 = 0.0717$ $\theta_2 = 0.2613$ $\theta_3 = 0.0224$	$\theta_1 = 0.0717$ $\theta_2 = 0.2613$ $\theta_3 = 0.0224$
$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3\sqrt{X_t}dW_t$	$\theta_1 = 0.0732$ $\theta_2 = 0.1452$ $\theta_3 = 0.0652$	$\theta_1 = 0.0722$ $\theta_2 = 0.2184$ $\theta_3 = 0.0667$	$\theta_1 = 0.0721$ $\theta_2 = 0.2189$ $\theta_3 = 0.0667$	$\theta_1 = 0.0721$ $\theta_2 = 0.2190$ $\theta_3 = 0.0667$	$\theta_1 = 0.0721$ $\theta_2 = 0.2189$ $\theta_3 = 0.0667$
$dX_t = X_t[\theta_2 - (\theta_3^2 - \theta_2\theta_1)X_t]dt + \theta_3 X_t^{3/2} dW_t$	$\theta_1 = 15.0190$ $\theta_2 = 0.1771$ $\theta_3 = 0.8059$	$\theta_1 = 15.1565$ $\theta_2 = 0.1813$ $\theta_3 = 0.8211$	$\theta_1 = 15.1411$ $\theta_2 = 0.1821$ $\theta_3 = 0.8211$	$\theta_1 = 15.1413$ $\theta_2 = 0.1821$ $\theta_3 = 0.8211$	$\theta_1 = 15.1414$ $\theta_2 = 0.1821$ $\theta_3 = 0.8211$
$dX_t = \theta_2(\theta_1 - X_t)dt + \theta_3 X_t^{\theta_4} dW_t$	$\theta_1 = 0.0808$ $\theta_2 = 0.0971$ $\theta_3 = 0.7224$ $\theta_4 = 1.4607$		$\theta_1 = 0.0842$ $\theta_2 = 0.0886$ $\theta_3 = 0.7792$ $\theta_4 = 1.4812$	$\theta_1 = 0.0842$ $\theta_2 = 0.0886$ $\theta_3 = 0.7792$ $\theta_4 = 1.4812$	

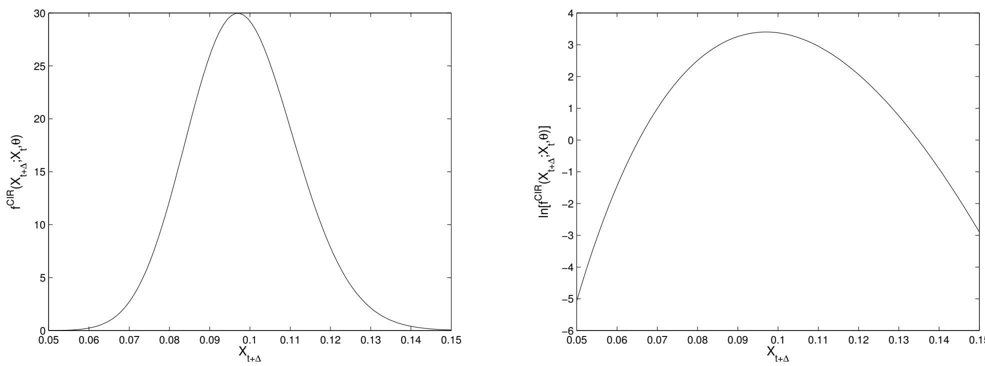

(a) Transition Density (b) Log Transition Density

Figure 4.2: Conditional Transition Density for CIR Process

Based on the parameters $\theta = (.06, .5, .15)$, $\Delta = 1/12$, and $X_t = .1$.

4.3.2 Testing with the CIR process

To further examine the properties of the new approximation for the conditional transition density we consider various benchmarks using the CIR process as described in Section 4.3.1. These tests are similar to those presented in Ait-Sahalia [2002] and Durham and Gallant [2002]. In order to stay consistent with previous benchmarking in this area we use base case parameter values of $\theta^0 = (.06, .5, .15)$, $\Delta = 1/12$, and $X_0 = .1$ in our experiments. For these parameter values the true conditional transition density and its log are presented in Figure 4.2.

CIR Benchmarks

Both Ait-Sahalia [2002] and Durham and Gallant [2002] develop a number of benchmarks with which one may test the efficiency of new estimation techniques using the CIR model. We follow the precedence they have set utilizing three benchmarks.

First we present a visual analysis of the approximation errors associated with the various techniques discussed in this paper. That is for a fixed point, $X_t = X_0$, the approximate conditional transition density for a series of points, $X_{t+\Delta} \in [0.05, 0.15]$, is computed and compared with the true conditional transition density. Since the end goal is to accurately approximate the log-likelihood function we consider the errors in approximating the log of the density as opposed to its level. As pointed out by Durham and Gallant [2002], the errors associated with approximating the level of the density do not provide an intuitive way of analyzing the effect such methods will have on constructing the log-likelihood function.

In order to better understand the ability of the proposed approximations in computing the log-likelihood function, a second benchmark is used. For this benchmark a series of $L = 10,000$ observations are simulated using draws from the non-central chi-square distribution, given a starting value $X_t = X_0$. The true and approximate conditional transition densities are then computed for each observation generated. Two methods for analyzing the accuracy of the approximation are considered. The first is the root mean squared error (RMSE),

$$RMSE = \left\{ \frac{1}{L} \sum_{i=0}^{L-1} \left(\ln(\hat{f}(X_{i+1}; X_i, \theta^0)) - \ln(f^{CIR}(X_{i+1}; X_i, \theta^0)) \right)^2 \right\}^{\frac{1}{2}},$$

where $\hat{f}$ represents the approximation of the density and f^{CIR} is the true conditional transition density as defined in (4.20). The second method is the maximum approximation error relative to the maximum value of the conditional transition density (MRE),

$$MRE = \max_i \left\{ \frac{\left| \ln(\hat{f}(X_{i+1}; X_i, \theta^0)) - \ln(f^{CIR}(X_{i+1}; X_i, \theta^0)) \right|}{\max_i [\ln(f^{CIR}(X_{i+1}; X_i, \theta^0))]} \right\}.$$

Using this benchmark one is also able to gauge the computational efficiency of the various approximations techniques.

The third benchmark employs Monte Carlo experiments in order to examine the ability of the approximation techniques to produce accurate parameter estimates in a manner more rigorous than in the previous section. For the experiment a series of $N = 512$ data sets with $L = 1,000$ observations are generated. Then the parameter estimates are computed using both the exact log-likelihood function and it approximations. For the exact

MLE we compute the error in estimating the true parameters using the mean error (ME),

$$ME_{TRUE-MLE} = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{\theta}_{MLE} - \theta^0\right).$$

The estimates obtained from using the approximation techniques are then compared to the exact MLE using the mean error,

$$ME_{MLE-APPROX} = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{\theta}_{MLE} - \hat{\theta}_{APPROX}\right), \tag{4.22}$$

for a given approximation technique denoted by APPROX. Durham and Gallant [2002] suggest that in order to be considered accurate an error of 1% of the error associated with the MLE should be achievable using the approximation.

CIR Benchmark Results

Plots of the approximation errors using the Euler method with the untransformed process, the Euler method with the transformed process, and the second order Milstein approximation are presented in Figures 4.3, 4.4, and 4.5 respectively. In each of these tests four approximations are considered. The methods are denoted by $f^{(I)}$, where I represents the number of intervals used to compute the conditional transition density. In all cases we use a quadrature order of $M = 15$. Note that in the case of Richardson extrapolation the method is denoted as the simple sum of the density estimates, excluding the necessary coefficients. This approach is simply to reduce the complexity of the notation and will be used throughout the remainder of this paper.

It is strikingly apparent that the constant diffusion transformation provides a significant improvement for the Euler method by an order of magnitude. Another order of magnitude improvement may be gained by implementing a higher order approximation such as the Milstein approach. The technique of Shoji & Ozaki provides a similar benefit to the of Milstein. Furthermore, it may be seen that the use of numerical quadrature techniques to increase the number of sub-densities used in the approximation, does indeed reduce the errors associated with the approximation. Most notable however, is the improvement made

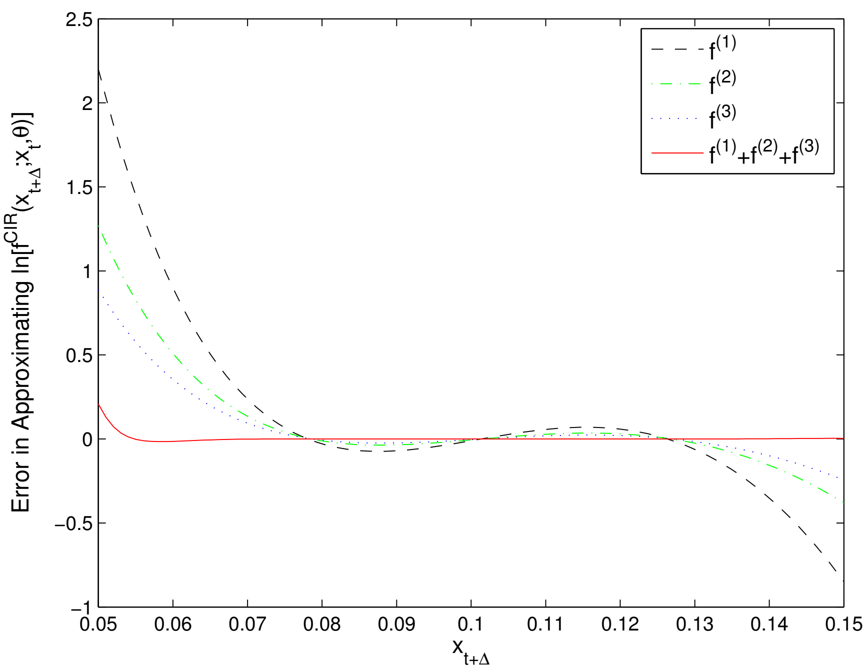

Figure 4.3: Approximation Error with Untransformed Euler

once the various approximations are combined through Richardson extrapolation.

These points are further illustrated in Table 4.2 which presents the results of the second set of benchmarks analyzing the ability of the approximation to compute the conditional transition densities for a series of $L = 10,000$ simulated observations. Again it may be seen that the constant diffusion transformation provides an order of magnitude improvement in terms of reducing the RMSE. While a further order of magnitude improvement is possible by implementing one of the higher order approximations such as the second order Milstein or Shoji & Ozaki specification. Such alternatives significantly improve the approximation without a substantial increase in computation time.[6]

Further reduction in the approximation error may be obtained through incorporating sub-densities into the transition density. Though again the true benefit of being able to reduce the step size of the approximation comes from the ability to apply Richardson extrapolation, which takes very little additional computation. Using the Shoji & Ozaki

[6]All run times were computed on an Intel Core Duo 2.0 Ghz processor with 2 GB of RAM, running OS X 10.5.6 and MATLAB 2007a.

Table 4.2: Approximation Errors for the Log Conditional Transition Density

Moment Approximation	**Density Approximation**	**M**	**RMSE**	**MRE**	**Time (sec.)**
Untransformed Euler	$f^{(1)}$		0.14499	1.10943	0.01
	$f^{(2)}$	10	0.07009	0.38198	0.02
	$f^{(3)}$	10	0.04716	0.22640	0.32
	$f^{(1)} + f^{(2)}$	10	0.02757	0.27752	0.03
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00681	0.06355	0.36
True Moments	$f^{(1)}$		0.14267	0.91936	0.01
	$f^{(2)}$	10	0.07023	0.34593	0.03
	$f^{(3)}$	10	0.04743	0.20962	0.46
	$f^{(1)} + f^{(2)}$	10	0.02550	0.27256	0.04
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00647	0.07487	0.52
Transformed Euler	$f^{(1)}$		0.03649	0.12471	0.01
	$f^{(2)}$	10	0.01832	0.05906	0.03
	$f^{(3)}$	10	0.01223	0.03858	0.32
	$f^{(1)} + f^{(2)}$	10	0.00275	0.04071	0.03
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00035	0.00386	0.34
	$f^{(1)} + f^{(2)} + f^{(3)}$	5	0.01502	0.03198	0.11
	$f^{(1)} + f^{(2)} + f^{(3)}$	25	0.00034	0.00334	0.75
	$f^{(1)} + f^{(2)} + f^{(3)}$	25	0.00033	0.00335	2.02
Milstein	$f^{(1)}$		0.00592	0.05878	0.01
	$f^{(2)}$	10	0.00129	0.01171	0.05
	$f^{(3)}$	10	0.00056	0.00482	0.67
	$f^{(1)} + f^{(2)}$	10	0.00036	0.00309	0.05
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00005	0.00062	0.71
	$f^{(1)} + f^{(2)} + f^{(3)}$	5	0.00460	0.00634	0.20
	$f^{(1)} + f^{(2)} + f^{(3)}$	15	0.00004	0.00044	1.56
	$f^{(1)} + f^{(2)} + f^{(3)}$	25	0.00004	0.00044	4.22
Shoji & Ozaki	$f^{(1)}$		0.00519	0.04907	0.01
	$f^{(2)}$	10	0.00121	0.01019	0.07
	$f^{(3)}$	10	0.00053	0.00424	1.25
	$f^{(1)} + f^{(2)}$	10	0.00026	0.00178	0.08
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00003	0.00021	1.32
	$f^{(1)} + f^{(2)} + f^{(3)}$	5	0.00461	0.00634	0.37
	$f^{(1)} + f^{(2)} + f^{(3)}$	15	0.00002	0.00021	2.92
	$f^{(1)} + f^{(2)} + f^{(3)}$	25	0.00002	0.00025	7.94

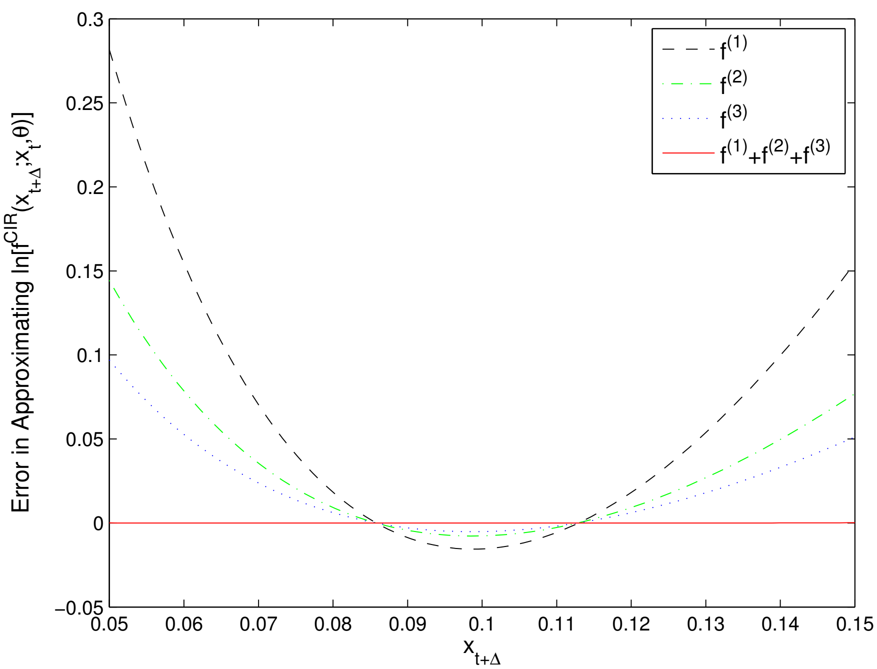

Figure 4.4: Approximation Error with Transformed Euler

approximation for the moments of the Gaussian sub-density and combining the versions with one, two, and three subintervals results in an MRE of 0.00021. Or in other words, the maximum error is only one fiftieth of one percent of the maximum value of the true log conditional transition density. To obtain this result the practitioner needs not approximate an integral higher than two dimensions, which may be computed relatively quickly using numerical quadrature methods. We note that the second order Milsten approximation provides nearly the same level of accuracy as the Shoji & Ozaki specification. However, the lack of exponential operators in the Milstein approximation allows for quicker computation.

We also consider the the use of the true moments for the CIR process in place of the approximations presented in Section 4.1. For the CIR process the true conditional moments are

$$E[X_{t+\Delta}\,|X_t,\theta] = \theta_1 + e^{-\theta_2\Delta}(X_t - \theta_1)$$

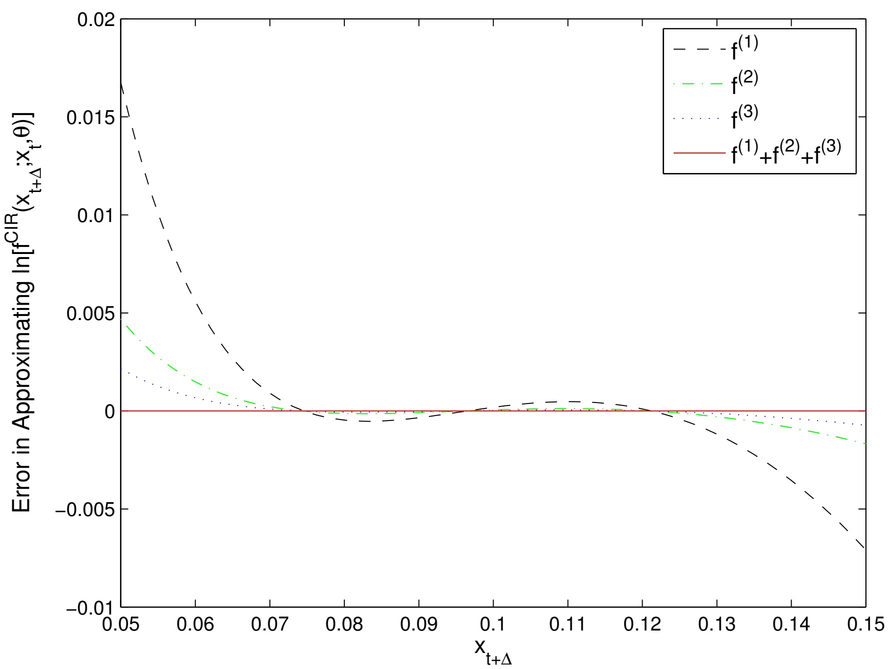

Figure 4.5: Approximation Error with Second Order Milstein Approximation

and

$$Var[X_{t+\Delta}\,|X_t,\theta] = \frac{\theta_3^2\left(1-e^{-\theta_2\Delta}\right)}{\theta_2}\left[\theta_1\left(1-e^{-\theta_2\Delta}\right)+X_te^{-\theta_2\Delta}\right].$$

However, it may be seen from the results in Table 4.2 that the use of the true conditional moments provides very little increase in accuracy over the first order Euler approximation for the untransformed process. This result follows with the discussion of Section 4.1.2 which notes that the constant diffusion transformation is required to attain higher order approximations for the conditional transition density. It is worth noting that for many cases, including the CIR process, the SDE of the transformed process does not have a linear drift function, which for the most part rules out the possibility of deriving closed form expressions for the true conditional moments.

In addition to the choice of moment approximation and the number of sub-intervals, a major determinant of the computational complexity and accuracy of the approximation is the order of the numerical quadrature. As may be seen in Table 4.2, up to a point increasing the accuracy of the numerical integral will result in a reduction of the approximation error.

For example moving from $M = 5$ to $M = 10$. This result is consistent across the moment specifications. However, when increasing the quadrature order a point will be reached where the error of the numerical integral is dominated by the error associated with the Gaussian assumption and moment approximations. In this case a further increase in the order of the numerical integral does not result in a significant improvement of the conditional transition density approximation, only an increase in the computation time. For example increasing the order from $M = 15$ to $M = 25$ in the case of $f^{(1)} + f^{(2)} + f^{(3)}$ using the Shoji & Ozaki specification reduces the RMSE by less than 10^{-6}, though increases the computation time by 172%.

Given the importance of quadrature order to both accuracy and efficiency we explore this point further through a simulation based analysis of the order of the approximation error. As discussed in Section 4.2 for the case of a relatively small integration error the relationship

$$f^{(I)}(X_{t+\Delta}; X_t, \theta) = f(X_{t+\Delta}; X_t, \theta) + a\left(\frac{\Delta}{I}\right)^n + O(\Delta^{n+1}). \tag{4.23}$$

will hold, where n is the order of the approximation error. When applying the technique of Richardson extrapolation it is necessary for this condition to hold in order to be able to derive the correct weights. To estimate the order of the approximation error we simulated $N = 1,000$ observations of the CIR process and estimate the log mean absolute error version of (4.23),

$$\ln\left(\frac{1}{N}\sum_{n=1}^{N}\left|f^{(I)} - f\right|\right) = \ln(\hat{a}) + \hat{n}\ln\left(\frac{\Delta}{I}\right), \tag{4.24}$$

over the estimators $f^{(I)}$ $i = 1, ..., 4$.

Table 4.3 presents the estimated order of the approximation error for the quadrature order of $M = 10$. It is clear that the error is well described by (4.24). As expected the Euler approximation is of order one independent of the constant diffusion transformation, however the transformation does significantly reduce the error through its effect on the constant a. This is line with the results of the second benchmark presented in Table 4.2. Also as expected, the Milstein and Shoji & Ozaki approximations are of order two. These

Table 4.3: Mean Absolute Error for M=10

The error was computed for $N = 1,000$ simulated observations using the true parameter set $\theta^0 = (.06, .5, .15)$, step size $\Delta = 1/12$, and $X_t = .1$.

Moment Approximation	$\hat{a}$	$\hat{n}$	**Residual**
Euler	11.13	0.99	0.00
Euler (Transformed)	3.12	1.00	0.00
Milstein	0.89	2.00	0.01
Shoji & Ozaki	0.91	2.01	0.04
True Moments	10.38	0.99	0.00

Table 4.4: Mean Absolute Error for Milstein Approximation

The error was computed for $N = 1,000$ simulated observations using the true parameter set $\theta^0 = (.06, .5, .15)$, step size $\Delta = 1/12$, and $X_t = .1$.

M	$\hat{a}$	$\hat{n}$	**Residual**
4	0.00	-1.21	2.20
6	0.18	1.08	1.70
8	1.27	2.06	0.50
10	0.89	2.00	0.01

results also further confirm that the true conditional moments offer only slightly better performance the first order Euler approximation for the untransformed process. Plots of the error versus the time step Δ/I are presented in log scale in Figure 4.6.

These results suggest that (4.23) holds, indicating that the integration error is relatively small. In cases where integration over the sub-intervals is inaccurate the approximation of the conditional transition density will be inaccurate as well. Furthermore Richardson extrapolation will no longer be a viable technique for reducing the order of the approximation error. Table 4.4 presents estimates of (4.24) for the Milstein approximation over the quadrature errors $M \in 4, 6, 8, 10$. It is clear that the relationship in (4.23) is sensitive to the quadrature order, a point that is further illustrated in Figure 4.7 which plots the error versus the time step Δ/I in log scale. It appears that $M = 10$ is sufficient to provide an approximation of order Δ^2 for the Milstein approximation. This is in line with the results from the second benchmark in which increasing the order of approximation beyond $M = 10$ providing little benefit, as the integration error is dominated by that associated with the sub-density approximation.

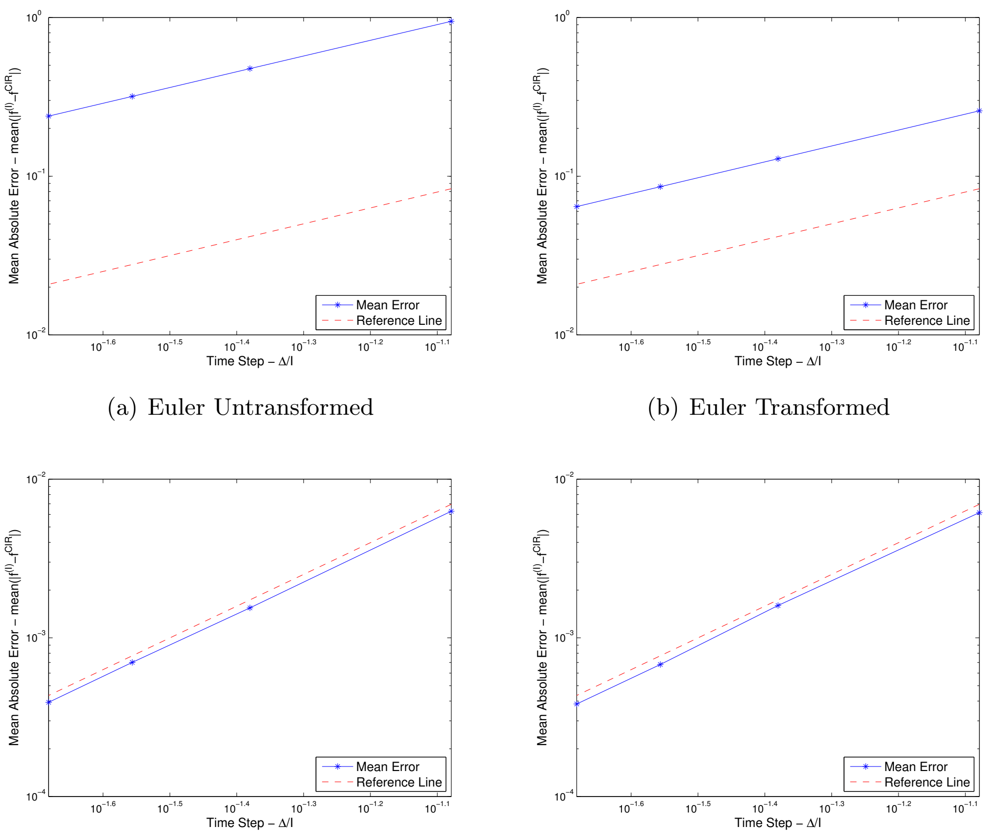

(a) Euler Untransformed

(b) Euler Transformed

(c) Milstein

(d) Shoji & Ozaki

Figure 4.6: Mean Absolute Error in Log Scale for M=10

The reference line has a slope of two for the Milstein and Shoji & Ozaki cases and a slope of one for the other two cases. The error was computed for $N = 1,000$ simulated observations using the true parameter set $\theta^0 = (.06, .5, .15)$, step size $\Delta = 1/12$, and $X_t = .1$.

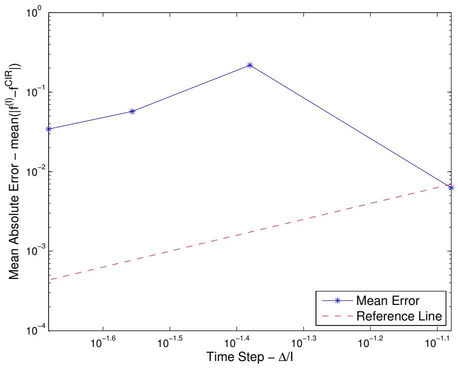

(a) M=4

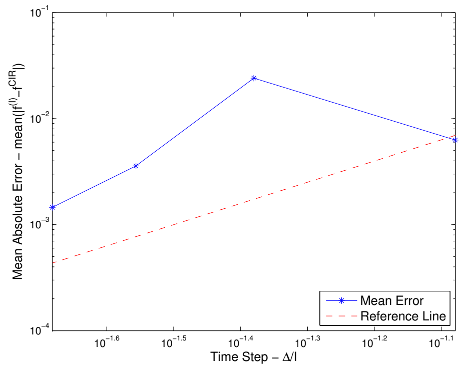

(b) M=6

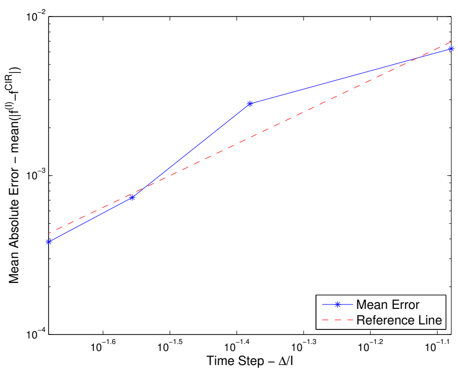

(c) M=8

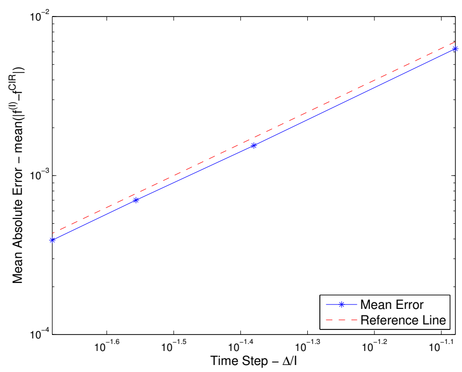

(d) M=10

Figure 4.7: Mean Absolute Error in Log Scale for Milstein Approximation

The reference line has a slope of two in all cases. The error was computed for $N = 1,000$ simulated observations using the true parameter set $\theta^0 = (.06, .5, .15)$, step size $\Delta = 1/12$, and $X_t = .1$.

In order to examine the efficiency of the purposed approximation in estimating the parameters of SDEs we turn to the third benchmark. Monte Carlo experiments are conducted in order to determine the accuracy of parameter estimates obtained used the approximate conditional transition density relative to the true density. The results for the experiments are presented in Table 4.5. As may be seen, the criteria suggested by Durham and Gallant [2002] of 1% of the error associated with the MLE is easily reached by the $f^{(1)} + f^{(2)}$ approximation. In fact the mean error across the simulations is below 10^{-5} for all three parameters in the case of $M = 15$. This requires only computing a one dimensional integral for each pair of observations. We note that Durham and Gallant [2002] achieved similar accuracy with the SMLE, however for this result they required sixteen subintervals. As a result a fifteen dimensional integral needed to be computed for each pair of observations, using Monte Carlo techniques.

It is worth noting that additional accuracy may be obtained by both increasing the order of the integral approximation and the number of sub-intervals included in the density approximation. An increase in the order of the numerical quadrature technique, M, will provide some improvement in the accuracy of the estimates up until the quadrature error becomes dominated by the density approximation error. As may be seen in Table 4.5, increasing the order from $M = 10$ to $M = 15$ for the case of $f^{(2)}$ does little to improve the accuracy of the estimates. However, it should be noted that in the case of $f^{(1)} + f^{(2)}$ the increase in order has a more significant effect. As discussed in Section 4.2 the ability of Richardson extrapolation to increase the order of the approximation is dependent upon the integration error being relatively small. From our testing it appears that through the use of Gaussian quadrature such accuracy in the numerical integration is easily achieved.

4.3.3 Additional Examples

To further examine the accuracy of this new approximation technique for a wider variety of SDEs we consider three additional examples. These models cover cases with non-linearities in both the drift and diffusion, strong skewness in the conditional transition density, and an inability to apply a constant diffusion transformation. In all of these complex situations the new method presented in this paper is able to provide computationally

Table 4.5: Mean Approximation Errors for Parameter Estimates

Following the benchmark of Durham and Gallant [2002] the mean error was computed for 512 experiments with $N = 1,000$ simulated observations using the true parameter set $\theta^0 = (.06, .5, .15)$, step size $\Delta = 1/12$, and $X_t = .1$. The error for the MLE is computed from the true parameters, whereas the error for the conditional transition density approximations is taken from the MLE. The standard deviations for the mean approximation errors are presented in parentheses.

Moment Approximation	**Density Approximation**	**M**	θ_1	θ_2	θ_3
Transformed Euler	$f^{(CIR)}$		0.00041 (0.00822)	0.04534 (0.11995)	0.00023 (0.00350)
	$f^{(1)}$		0.00026 (0.00003)	0.01780 (0.01139)	0.00366 (0.00075)
Milstein	$f^{(1)}$		-0.00000 (0.00003)	0.00185 (0.00834)	-0.00009 (0.00008)
	$f^{(2)}$	10	-0.00000 (0.00001)	0.00028 (0.00253)	-0.00002 (0.00002)
	$f^{(2)}$	15	-0.00000 (0.00001)	0.00021 (0.00147)	-0.00002 (0.00001)
	$f^{(3)}$	10	0.00000 (0.00000)	0.00008 (0.00063)	-0.00001 (0.00001)
	$f^{(1)} + f^{(2)}$	10	0.00000 (0.00001)	0.00017 (0.00263)	-0.00000 (0.00001)
	$f^{(1)} + f^{(2)}$	15	0.00000 (0.00000)	0.00009 (0.00159)	-0.00000 (0.00001)
	$f^{(1)} + f^{(2)} + f^{(3)}$	10	0.00000 (0.00000)	0.00001 (0.00071)	-0.00000 (0.00001)

accurate and efficient approximations of the conditional transition densities.

Duffie 2001

The model presented in Duffie [2001] provides a more complex example with highly non-linear drift and diffusion functions. The SDE has the form

$$dX_t = [\theta_1 + \theta_2 X_t + \theta_3 X_t \ln(X_t)]\, dt + (\theta_4 + \theta_5 X_t)^{\theta_6} dW_t.$$

It is worth noting that the Ornstein-Uhlenbeck, CIR, and CEV processes are all specific cases of this more general model. For the constant diffusion case the transformation has the form

$$Y_t = \frac{(\theta_4 + \theta_5 X_t)^{1-\theta_6}}{\theta_5(1-\theta_6)},$$

where, by Ito's Lemma, the transformed SDE will be

$$\begin{aligned} dY_t = & \left\{ \tfrac{-\theta_5\theta_6}{2\gamma(Y_t,\theta)} + \tfrac{[\theta_2+\theta_3\eta(Y_t,\theta)]\gamma(Y_t,\theta)}{\theta_5} + \left[\theta_1 - \tfrac{\theta_4\theta_2}{\theta_5} - \tfrac{\theta_3\theta_4}{\theta_5}\eta(Y_t,\theta)\right] [\gamma(Y_t,\theta)]^{\frac{-\theta_6}{1-\theta_6}} \right\} dt \\ & + dW_t, \end{aligned}$$

where

$$\gamma(Y_t,\theta) = -\theta_5(\theta_6 - 1)Y_t$$

and

$$\eta(Y_t,\theta) = \ln\left\{ \frac{[\gamma(Y_t,\theta)]^{\frac{1}{1-\theta_6}} - \theta_4}{\theta_5} \right\}.$$

Unfortunately the true conditional transition density for this process is unknown in closed form. However, an extremely accurate approximation may be obtained as the solution to the forward Kolmogorov (Fokker-Planck) equation

$$\begin{aligned} \frac{\partial f(X_{t+\Delta};X_t,\theta)}{\partial \Delta} = & \left\{ bb'' + (b')^2 - a' \right\} f(X_{t+\Delta};X_t,\theta) + [2bb' - a] \frac{\partial f(X_{t+\Delta};X_t,\theta)}{\partial X_{t+\Delta}} \\ & + \frac{b^2}{2}\frac{\partial^2 f(X_{t+\Delta};X_t,\theta)}{\partial X_{t+\Delta}^2}, \end{aligned} \tag{4.25}$$

subject to the initial condition that $f(X_t;\theta)$ is equal to the Dirac function centered at X_t. The dependence of the drift and diffusion functions on X_t and θ has been left out in order

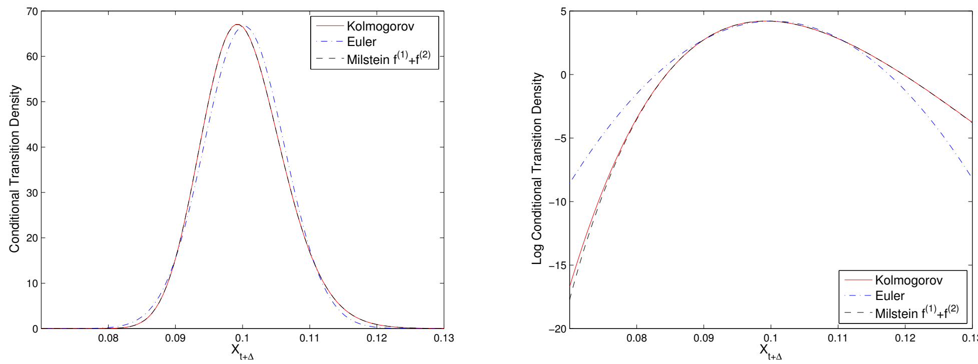

Figure 4.8: Conditional Transition Density Approximation for Duffie 2001

to simplify the notation. We solve the forward Kolmogorov equation using adaptive finite difference methods in time and 500 nodes in the space dimension. Though this method is computational inefficient for approximating the conditional transition density for the purpose of parameter estimation, it does provide an accurate description for the purpose of testing the new method.

For the purpose of testing we use the parameters $\theta = (0.03, 0.81, 0.47, 0.75, 0.76, 20.28)^T$ which are calibrated to Federal Funds rate data used in Section 4.3.1, along with $\Delta = 1/12$ and $X_t = 0.10$. Figure 4.8 presents the approximations for both the conditional transition density and its natural log. As may be seen, the first order Euler approximation on the untransformed process performs poorly relative to the new method using the Milstein approximation and $f^{(1)} + f^{(2)}$ with a quadrature order of $M = 10$.

Ait-Sahalia 1996

Another interesting example to consider is the process introduced by Ait-Sahalia [1996] where the SDE has the form

$$dX_t = \left(\theta_1 + \theta_2 X_t + \theta_3 X_t^2 + \frac{\theta_4}{X_t}\right) dt + \sqrt{\theta_5 + \theta_6 X_t + \theta_7 X_t^{\theta_8}}\, dW_t.$$

This process is of particular interest due both its complexity and the fact that there does not exist a closed form constant diffusion transformation. Despite the inability to

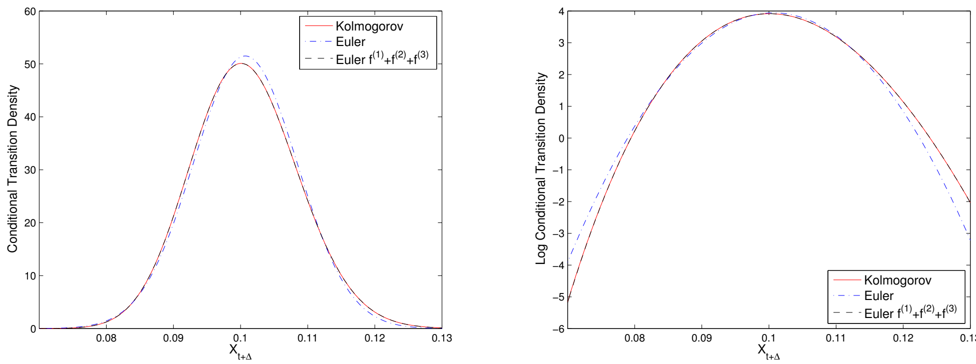

Figure 4.9: Conditional Transition Density Approximation for Ait-Sahalia 1996

obtain a constant diffusion version of the process the new method is still able to accurately approximate the conditional transition density. Again the true conditional transition density is not known is closed form, though we are able to accurately approximate it through a numerical solution to (4.25). As with the Duffie 2001 example we use paraemters calibrated from the Federal Funds rate data utilized in Section 4.3.1, such that $\theta = (-0.29, 3.36, -10.26, 0.0065, -0.00028, 0.010, 0.40, 37.76)^T$, along with $\Delta = 1/12$ and $X_t = 0.10$.

Since a constant diffusion transformation is unavailable for this process we apply the first order Euler approximation in order to obtain the conditional mean and variance for the Gaussian transition density. In order to obtain a third order approximation for the conditional transition density using the untransformed process we utilize $f^{(1)} + f^{(2)} + f^{(3)}$ with a quadrature order of $M = 10$. The approximation for the conditional transition density is presented in Figure 4.9 along with the standard Euler approximation and that obtained from solving the forward Kolmogorov equation. As is clear, the third order approximation obtained by using Richardson extrapolation to combine $f^{(1)}$, $f^{(2)}$, and $f^{(3)}$ is significantly more accurate than the standard first order Euler approximation. This suggests the method presented in this paper may be used to compute accurate approximations of the conditional transition density even for processes that may not be transformed to a constant diffusion case.

Power CIR

The last example we consider is that of the power CIR process, Y_t, where $Y_t = (X_t/\theta_1)^{\theta_4}$ and X_t follows the CIR SDE in (4.19). The conditional transition density for this process is strongly skewed for large values of θ_4. This model therefore provides an ideal way of testing the ability of the new approximation method to handle such skewness. The SDE will have the form

$$dY_t = \frac{\theta_4}{2\theta_1^{\theta_4}} \left(\theta_1 Y_t^{\frac{1}{\theta_4}}\right)^{\theta_4 - 1} \left[\theta_3^2(\theta_4 - 1) - 2\theta_1\theta_2 \left(Y_t^{\frac{1}{\theta_4}} - 1\right)\right] dt + \frac{\theta_3\theta_4}{\theta_1^{\theta_4}} \left(\theta_1 Y_t^{\frac{1}{\theta_4}}\right)^{\theta_4 - \frac{1}{2}} dW_t,$$

where the true conditional transition density is known to be

$$f(Y_{t+\Delta}; Y_t, \theta) = \frac{\theta_1^{\theta_4}}{\theta_4} Y_t^{\frac{1}{\theta_4} - 1} f^{CIR}(X_{t+\Delta}; X_t, \theta),$$

where $f^{CIR}(X_{t+\Delta}; X_t, \theta)$ is described by (4.20). The constant diffusion transformation has the form $Z_t = 2\theta_4 Y_t^{1/2\theta_4}$ where by Ito's Lemma

$$dZ_t = \left[\theta_4^2 \left(2\theta_2 - \frac{\theta_3^2}{2\theta_1}\right) \frac{1}{Z_t} - \frac{\theta_2}{2} Z_t\right] dt + \frac{\theta_3\theta_4}{\sqrt{\theta_1}} dW_t.$$

For testing purposes we use the same parameters as Ait-Sahalia [2002] where $\theta = (0.08, 0.50, \sqrt{0.02}, 7)^T$, $\Delta = 1/12$, and $Y_t = 4.77$. Figure 4.10 presents the true conditional transition density along with the first order Euler approximation using the untransformed process and the approximation using Milstein and $f^{(1)} + f^{(2)}$ with $M = 10$. Despite the strong skewness in the conditional transition density the new approximation method is capable of accurately computing the density. To further illustrate the accuracy of the new approximation Figure 4.11 presents the approximation error.

4.4 Concluding Remarks

Given the popularity of models utilizing continuous time stochastic variables, estimation of the unknown parameters in SDEs using discretely observed data is important in the field of economics as well as finance. The lack of explicit conditional transition densities

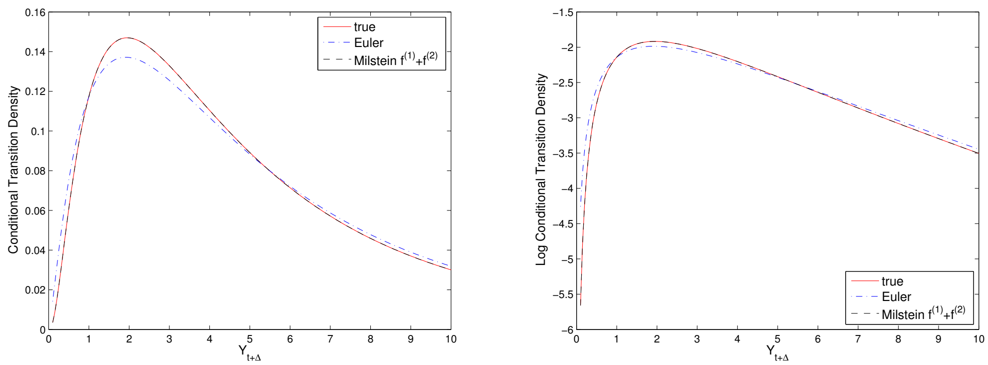

Figure 4.10: Conditional Transition Density Approximation for Power CIR

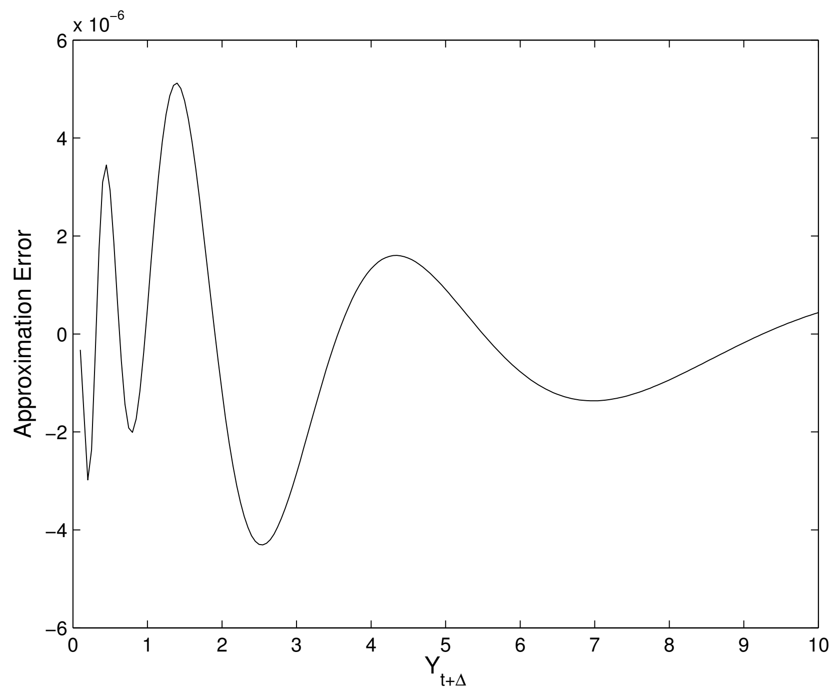

Figure 4.11: Conditional Transition Density Approximation Error for Power CIR

for all but a few SDEs eliminates the possibility of obtaining the true MLE in most applications. We have presented an approximation to the conditional transition density that produces a QMLE which is within a negligible distance of the MLE. In tests with the CIR process when using the new approximation for the conditional transition density, parameter estimates that are within 1% of the error associated with the MLE are easily obtained. This is accomplished through the use of Gauss Hermite quadrature for the computation of a one dimensional integral in the density approximation. As such this new technique has a significant computational advantage over similar SMLE methods. In addition to its computational accuracy and efficiency the approximation is very easy to implement, making it a desirable alternative to more complicated QMLE methods such as those utilizing Hermite expansions of the conditional transition density.

Given the accuracy of Gauss Hermite quadrature the conditional transition density obtained from the Chapman-Kolmogorov equation may be computed quickly with very few nodes. The results presented suggest that using only 10 nodes in the one dimensional quadrature rules is adequate to ensure that the integration error is of an order less than that of the sub-density approximation. For the case of introducing two sub-intervals with a second order sub-density approximation, such as that of Milstein, the $f^{(3)}$ approximation to the conditional transition density may be computed with only 100 nodes. As a result the log-likelihood for a sample of 1,000 observations may be easily computed in under a tenth of a second. Using the $f^{(2)}$ approximation with moments derived from the Milstein approximation, a case requiring only one dimensional integration with 10 nodes, the log-likelihood for 1,000 observations is computed in around one one hundredth of a second. This represents a negligible increase in computation time from the case where no sub-intervals are used but a significant increase in accuracy.

Furthermore, we find the use of Richardson extrapolation to be highly beneficial in producing accurate approximations of the conditional transition density. Through its use higher order approximations may be obtained with nearly no additional computational burden. The second order conditional transition density approximation $f^{(3)}$ that is obtained by using moments derived from the Milstein approximation, may be easily expanded to a fourth order approximation by combining with the $f^{(1)}$ and $f^{(2)}$ densities. In this case the log-likelihood for a sample of 1,000 observations may still be computed in under one tenth of a second. However, our testing suggests that for monthly observations the third order approximation $f^{(1)} + f^{(2)}$ with moments derived from the Milstein approximation is more than adequate. This requires only one dimensional integration with 10 nodes allowing for the log-likelihood of 1,000 observations to be computed in around one one hundredth of a second. However, the ability to utilize Richardson extrapolation to increase the approximation order in an efficient manner is dependent upon quick and accurate numerical integration of the Chapman-Kolmogorov equation. For this purpose the use of Gauss Hermite quadrature, as suggested by this paper, is the ideal choice.

We also find that the constant diffusion transformation significantly reduces the error associated with the assumption of a Gaussian transition density. In addition the transformation allows for the use of higher order approximations through the implementation of Shoji & Ozaki's or Milstein's method for determining the moments of the Gaussian density. These techniques along with Richardson extrapolation allow for higher order approximations to be obtained with the inclusion of fewer sub-intervals. They are not required, however, to obtain accurate approximations of the conditional transition density. As demonstrated in Sections 4.3.2 and 4.3.3 accurate high order approximations may still be obtained with the use of the first order Euler approximation for the moments. This is useful for models in which a closed form constant diffusion transformation is unavailable.

BIBLIOGRAPHY

Y. Ait-Sahalia. Testing continuous-time models of the spot interest rate. *Review of Financial Studies*, pages 385–426, 1996.

Y. Ait-Sahalia. Transition densities for interest rate and other nonlinear diffusions. *Journal of Finance*, pages 1361–1395, 1999.

Y. Ait-Sahalia. Maximum Likelihood Estimation of Discretely Sampled Diffusions: A Closed-Form Approximation Approach. *Econometrica*, 70(1):223–262, 2002.

L.H.R. Alvarez and R. Stenbacka. Adoption of uncertain multi-stage technology projects: a real options approach. *Journal of Mathematical Economics*, 35(1):71–97, 2001.

G. Bakshi and N. Ju. A Refinement to Ait-Sahalia's (2002)" Maximum Likelihood Estimation of Discretely Sampled Diffusions: A Closed-Form Approximation Approach. *Journal of Business*, 78(5):2037–2036, 2005.

M. Balikcioglu. *Essays on Environmental and Computational Economics*. PhD thesis, 2008.

V. Bally and D. Talay. The law of the Euler scheme for stochastic differential equations: II. Convergence rate of the density. *Rapport De Recherche - Institut National De Recherche En Informatique Et En Automatique*, 1995.

C. Bartsch and E. Collaton. Coming clean for economic development: A resource book on environmental cleanup and economic development. Technical report, PB–96-130737/XAB, Northeast-Midwest Inst., Washington, DC (United States), 1995.

M.W. Brandt and P. Santa-Clara. Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets. *Journal of Financial Economics*, 63(2):161–210, 2002.

K.A. Brekke and B. Oksendal. Optimal Switching in an Economic Activity Under Uncertainty. *SIAM Journal on Control and Optimization*, 32(4):1021–1036, 1994.

M.J. Brennan and E.S. Schwartz. Evaluating Natural Resource Investments. *Journal of Business*, 58(2):135, 1985.

A. Cadenillas and F. Zapatero. Classical and impulse stochastic control of the exchange rate using interest rates and reserves. *Mathematical Finance*, 10(2):141–156, 2000.

D.R. Capozza and Y. Li. Optimal Land Development Decisions. *Journal of Urban Economics*, 51(1):123–142, 2002.

D.R. Capozza and G.A. Sick. Valuing long-term leases: The option to redevelop. *The Journal of Real Estate Finance and Economics*, 4(2):209–223, 1991.

P. Carr. The Valuation of Sequential Exchange Opportunities. *Journal of Finance*, 43(5): 1235–1256, 1988.

K.E. Case and R.J. Shiller. Is There a Bubble in the Housing Market? *Brookings Papers on Economic Activity*, 2(2003):299–342, 2003.

J.A. Chalmers and S.A. Roehr. Issues in the Valuation of Contaminated Property. *The Appraisal Journal*, 61(1):28–41, 1993.

KC Chan, G.A. Karolyi, F.A. Longstaff, and A.B. Sanders. An empirical comparison of alternative models of the short-term interest rate. *Journal of Finance*, pages 1209–1227, 1992.

J.P. Chancelier, B. 0ksendal, and A Sulem. Combined stochastic control and optimal stopping, and application to numerical approximation of combined stochastic and impulse control. *Stochastic Financial Mathematics, Steklov Math. Inst., Moscow*, 237:149–173, 2002.

M.H. Chang. Hereditary Portfolio Optimization with Taxes and Fixed Plus Proportional Transaction Costs—Part I. *Journal of Applied Mathematics and Stochastic Analysis*, 2007, 2007.

K.M. Chilton. The Myth of the 'Environmental Problem': Cleanup Costs and Brownfield Redevelopment. *Public Works Management and Policy*, 2(3):220–230, 1998.

E. Collaton and C. Bartsch. Industrial Site Reuse and Urban Redevelopment–An Overview. *Cityscape: A Journal of Policy Development and Research*, 2(3):17–61, 1996.

J.C. Cox, J.E. Ingersoll Jr, and S.A. Ross. A Theory of the Term Structure of Interest Rates. *Econometrica*, 53(2):385–408, 1985.

C.R. Cunningham. House price uncertainty, timing of development, and vacant land prices: Evidence for real options in Seattle. *Journal of Urban Economics*, 59(1):1–31, 2006.

T. De Luca, F. Facchinei, and C. Kanzow. A semismooth equation approach to the solution of nonlinear complementarity problems. *Mathematical Programming*, 75(3):407–439, 1996.

Avinash K. Dixit and Robert S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, 1994.

D. Duffie. *Dynamic asset pricing theory*. Princeton University Press Princeton, NJ, 2001.

Garland B. Durham and A. Ronald Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business and Economic Statistics*, 20(3):297–316, 2002.

O. Elerian, S. Chib, and N. Shephard. Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, pages 959–993, 2001.

B. Eraker. MCMC analysis of diffusion models with application to finance. *Journal of Business and Economic Statistics*, 19(2):177–191, 2001.

R. D. Espinoza and L. X. Luccioni. An approximate solution for perpetual american option with time to build: The value of environmental remediation investment projects. In *Real Options Annual International Conference*, 2005.

W. Feller. The parabolic differential equations and the associated semi-groups of transformations. *Annals of Mathematics*, pages 468–519, 1952.

A. Fischer. A special newton-type optimization method. *Optimization*, 24(3):269–284, 1992.

J.D. Fisher, G.H. Lentz, and K.S.M. Tse. Valuation of the Effects of Asbestos on Commercial Real Estate. *Journal of Real Estate Research*, 7:331–350, 1992.

D. Florens-Zmirou. Approximate discrete-time schemes for statistics of diffusion processes. *Statistics*, 20(4):547–557, 1989.

J.D. Hamilton. *Time Series Analysis*. Princeton University Press, 1994.

A.S. Holland, S.H. Ott, and T.J. Riddiough. The Role of Uncertainty in Investment: An Examination of Competing Investment Models Using Commercial Real Estate Data. *Real Estate Economics*, 28(1):33–64, 2000.

AS Hurn, JI Jeisman, and KA Lindsay. Seeing the wood for the trees: A critical evaluation of methods to estimate the parameters of stochastic differential equations. *Journal of Financial Econometrics*, 5(3):390, 2007.

C.S. Jones. Bayesian estimation of continuous-time finance models. *Unpublished paper, Simon School of Business, University of Rochester*, 1999.

K.L. Judd. *Numerical Methods in Economics*. MIT Press, 1998.

SE Kaiser. Commentary: Brownfield National Partnership. *Public Works Management and Policy*, 3(2):196–201, 1998.

CT Kelley. *Solving Nonlinear Equations with Newton's Method*. Society for Industrial Mathematics, 2003.

M. Kessler. Estimation of an ergodic diffusion from discrete observations. *Scandinavian Journal of Statistics*, 24:211–229, 1997.

P.E. Kloeden and E. Platen. *Numerical solution of stochastic differential equations*. Springer, 1992.

J. Lee and D.A. Paxson. Valuation of R&D real American sequential exchange options. *R&D Management*, 31(2):191–201, 2001.

G.H. Lentz and K.S.M. Tse. An option pricing approach to the valuation of real estate contaminated with hazardous materials. *The Journal of Real Estate Finance and Economics*, 10(2):121–144, 1995.

A.W. Lo. Maximum likelihood estimation of generalized Ito processes with discretely sampled data. *Econometric Theory*, pages 231–247, 1988.

R.R. Lumley and M. Zervos. A Model for Investments in the Natural Resource Industry with Switching Costs. *Mathematics of Operations Research*, 26(4):637–653, 2001.

V. Ly Vath, M. Mnif, and H. Pham. A model of optimal portfolio selection under liquidity risk and price impact. *Finance and Stochastics*, 11(1):51–90, 2007.

Saman Majd and Robert S. Pindyck. Time to build, option value, and investment decisions. *Journal of Financial Economics*, 18:7–27, 1987.

C.J. Mayer and C.T. Somerville. Land use regulation and new construction. *Regional Science and Urban Economics*, 30(6):639–662, 2000a.

C.J. Mayer and C.T. Somerville. Residential Construction: Using the Urban Growth Model to Estimate Housing Supply. *Journal of Urban Economics*, 48(1):85–109, 2000b.

P.B. Meyer and T.S. Lyons. Lessons from Private Sector Brownfield Redevelopers. *Journal of the American Planning Association*, 66(1):46–57, 2000.

P.B. Meyer and H.W. VanLandingham. Reclamation and Economic Regeneration of Brownfields. *Reviews of Economic Development Literature and Practice*, 1, 2000.

M.J. Miranda and P.L. Fackler. *Applied Computational Economics and Finance*. MIT Press, 2002.

G. Mundaca and B. Øksendal. Optimal stochastic intervention control with application to the exchange rate. *Journal of Mathematical Economics*, 29(2):225–243, 1998.

B. Mundy. The impact of hazardous and toxic material on property value: Revisited. *Appraisal Journal*, 60(4):463–471, 1992a.

B. Mundy. The Impact of Hazardous Materials on Property Value. *The Appraisal Journal*, 60(2):155–62, 1992b.

JP O'Brien. EPA's Landowner Liability Guidance. *Toxics Law Reporter*, 4(7):184–191, 1989.

B. Oksendal and B. Karsten. *Stochastic Differential Equations: An Introduction with Applications*. Springer, 1998.

B. Oksendal and A. Sulem. Optimal Consumption and Portfolio with Both Fixed and Proportional Transaction Costs. *SIAM Journal on Control and Optimization*, 40:1765, 2002.

B.K. Øksendal and A. Sulem. *Applied stochastic control of jump diffusions*. Springer Verlag, 2005.

S. Panayi and L. Trigeorgis. Multi-stage Real Options: The Cases of Information Technology Infrastructure and International Bank Expansion. *Quarterly Review of Economics and Finance*, 38:675–692, 1998.

P.J. Patchin. Valuation of Contaminated Properties. *The Appraisal Journal*, 56(1):7–16, 1988.

P.J. Patchin. Contaminated Properties—Stigma Revisited. *The Appraisal Journal*, 59(2): 167–172, 1991.

D.A. Paxson. Sequential American Exchange Property Options. *The Journal of Real Estate Finance and Economics*, 34(1):135–157, 2007.

A. R. Pedersen. A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scandinavian Journal of Statistics*, 22:55–71, 1995.

L. Quigg. Empirical Testing of Real Option-Pricing Models. *Journal of Finance*, 48(2): 621–640, 1993.

L. Quigg. Optimal Land Development. *Real Options in Capital Investment: Models, Strategies, and Applications*, pages 265–280, 1995.

T. Richards. Valuing contaminated land and property: theory and practice. *Journal of Property Valuation & Investment*, 14:4–17, 1996.

M.J. Rogers, A. Gupta, and C.D. Maranas. Real Options Based Analysis of Optimal Pharmaceutical Research and Development Portfolios. *Industrial and Engineering Chemistry Research*, 41(25):6607–6620, 2002.

E.S. Schwartz and L. Trigeorgis. *Real Options And Investment Under Uncertainty: Classical Readings and Recent Contributions*. MIT Press, 2004.

I. Shoji and T. Ozaki. Estimation for nonlinear stochastic differential equations by a local linearization method. *Stochastic Analysis and Applications*, 16:733–752, 1998.

R.A. Simons. How Many Urban Brownfields are Out There? An Economic Base Contraction Analysis of 31 US Cities. *Public Works Management & Policy*, 2(3):267–273, 1998.

M. Singer, L. Milligan, E. Stasiak, D. Borak, T. Groeneveld, D. Pickett, D. Clarke, and N. Mishkovsky. Brownfields Redevelopment: A Guidebook for Local Governments and Communities. *Washington, DC: ICMA*, 2001.

CF Sirmans. Research on discounted cash flow models. *Real Estate Finance*, 13(4):93–95, 1997.

D. Sun and L. Qi. On NCP-Functions. *Computational Optimization and Applications*, 13 (1):201–220, 1999.

P. Syms. Contaminated land and other forms of environmental impairment: an approach to valuation. *Journal of Property Valuation and Investment*, 14(2):38–47, 1996.

S. Tonin. 8. What is the Value ofBrownfields? AReview of Possible Approaches. *Valuing Complex Natural Resource Systems: The Case of the Lagoon of Venice*, 2006.

L. Trigeorgis. The Nature of Option Interactions and the Valuation of Investments with Multiple Real Options. *Journal of Financial and Quantitative Analysis*, 28(1):1–20, 1993a.

L. Trigeorgis. Real options and interactions with financial flexibility. *Financial Management*, 22(3):202–224, 1993b.

J.T. Williams. Real estate development as an option. *The Journal of Real Estate Finance and Economics*, 4(2):191–208, 1991.

A.R. Wilson. The Environmental Opinion: Basis for an Impaired Value Opinion. *The Appraisal Journal*, 62(3):410–23, 1994.

V.I. Zakamouline. European option pricing and hedging with both fixed and proportional transaction costs. *Journal of Economic Dynamics and Control*, 30(1):1–25, 2006.

APPENDIX

A Approximating the Development Stage Value Function

The value function for the development stage is defined by the variational inequality in (2.6)-(2.8) and the additional conditions (2.10) and (2.11). The specific form of the value function is unknown and therefore projection techniques as presented in Miranda and Fackler [2002] and Judd [1998] are used to approximate the function. Let

$$V^D(s,k) \approx \phi(s)c^D(k), \tag{A.1}$$

where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^q$ is a set of q basis functions for a family of approximating functions, and $c^D(k)$ is a q-dimensional vector of approximating coefficients when k units may still be constructed. In order to reduce the dimensionality of the approximation problem the change in value with respect to a change in k is approximated using a finite forward difference such that

$$\frac{\partial V^D}{\partial k} \approx \phi(s)\frac{c^D(k) - c^D(k-\Delta)}{\Delta}, \tag{A.2}$$

where Δ represents the step size. Given the approximations in (A.1) and (A.2) the variational inequality in (2.6)-(2.8) may be rewritten as

$$\rho\phi(s)c^D(k) \geq A(\theta^*)\left(\theta^* p - C_K\right) - C_M - A(\theta^*)\phi(s)\frac{c^D(k) - c^D(k-\Delta)}{\Delta} + L, \tag{A.3}$$

$$\rho\phi(s)c^D(k) \geq -C_S + L, \tag{A.4}$$

and

$$\phi(s)c^D(k) \geq \eta p k, \tag{A.5}$$

where L is the approximation of the differential generator

$$L = \left[\sum_{i=1}^{d} \mu_i(s)\frac{\partial \phi}{\partial s_i} + \frac{1}{2}\sum_{i=1}^{d}\sum_{j=1}^{d}\left[\sigma(s)\sigma(s)^T\right]_{ij}\frac{\partial^2 \phi}{\partial s_i \partial s_j}\right] c^D(k). \tag{A.6}$$

Given that one of the conditions in (A.3)-(A.5) must hold with equality the problem of obtaining the approximating coefficients, $c^D(k)$ for a given k may be represented by as a complementarity problem (CP) of the form

$$\min\left[B_1^D c^D(k) + b_1^D, B_2^D c^D(k) + b_2^D, B_3^D c^D(k) + b_3^D\right] = 0. \tag{A.7}$$

Given a vector of q nodal points $\hat{s}$, the CP is defined by

$$B_l^D = \begin{cases} \Phi & \text{if } l = 1 \\ \rho\Phi - \hat{L} & \text{if } l = 2 \\ \left[\rho + \frac{A(\theta^*)}{\Delta}\right]\Phi - \hat{L} & \text{if } l = 3 \end{cases},$$

and

$$b_l^D = \begin{cases} -\eta\hat{p}k & \text{if l=1} \\ C_S & \text{if } l = 2 \\ \frac{-A(\theta^*)}{\Delta}\Phi c^D(k-\Delta) - A(\theta^*)\left(\theta^* p - C_K\right) + C_M & \text{if } l = 3 \end{cases},$$

where Φ is the set of basis functions evaluated at the nodal points $\hat{s}$, $\hat{p}$ represents the price nodes within $\hat{s}$, $\hat{L}$ is the differential generator approximation at $\hat{s}$, and θ^* is the optimal choice for the relative price.

The optimal continuous control is defined by the numerical approximation to (2.11), such that

$$\theta^* = \frac{1}{2}\left(\frac{C_K + \Phi\frac{c^D(k) - c^D(k-\Delta)}{\Delta}}{p} - \frac{\kappa_0}{\kappa_1}\right). \tag{A.8}$$

The dependence of θ^* on the value function itself, implies that the CP in (A.7) is in fact an extended nonlinear complementarity problem (ENCP). To obtain the approximating coefficients, the Fischer-Burmeister approximation to the min function is applied in an iterative fashion in order to provide a semi-smooth system of nonlinear functions that may solved using a Newton-type method. This approach to solving regime switching models combined with a stochastic control is discussed in detail in Chapter 3.

Given the approximation of the boundary condition for $k = 0$ in (2.10)

$$\Phi c^D(0) = 0,$$

it may be seen that the vector of approximating coefficients, $c^D(0)$, at the end of the projects life is equal to a vector of zeros. Given this information we may iterate backwards in the K space, starting at $K = 0$, solving the ENCP in (A.7) for the next set of approximating coefficients at each Δ step.

B Approximating the Remediation Stage Value Function

As described in Section 2.3.2 the problem of obtaining an approximation for the value function while the investment is in the remediation and entitlement stage may be seen as the solution to set of $N+1$ problems of reduced complexity. This approach is based on the fact that after the regulation process has been concluded, $Y_t > 0$, the value function given a particular outcome may be obtained independent of the value under alternative outcomes. In the case of the outcome $Y_t = y$ where $y > 0$ and given the state (s, r) the value function is defined by the variational inequality

$$\rho V^R(s,r,y) \geq f(y,\beta) - \beta \frac{\partial V^R}{\partial r} + \mathcal{L}V^R(s,r,y) \quad \forall \beta \in \mathcal{B}, \tag{B.1}$$

and

$$V^R(s,r,y) \geq \eta p U(y)\epsilon - C_H r, \tag{B.2}$$

where $\mathcal{L}$ is the differential generator defined in (2.9). The approximation to the value function is defined as

$$V^R(s,r,y) \approx \phi(s) c^R(r,y), \tag{B.3}$$

where $\phi : \mathbb{R}^d \to \mathbb{R}^q$ is a set of q basis functions for a family of approximating functions, and $c^R(r,y)$ is a q-dimensional vector of approximating coefficients when r years of remediation remain and the outcome of the entitlement process is y. The dimensionality of the problem is reduced by approximating the change in value with respect to a change in r using a forward finite difference such that

$$\frac{\partial V^R}{\partial r} \approx \phi(s) \frac{c^R(r,y) - c^R(r-\Delta,y)}{\Delta}, \tag{B.4}$$

where Δ represents the step size. Given the approximations in (B.3) and (B.4) the variational inequality defining the value after the regulation process as defined in (B.1)-(B.2) may be rewritten as

$$\rho\phi(s)c^R(r,y) \geq f(y,\beta) - \beta\phi(s)\frac{c^R(r,y) - c^R(r-\Delta,y)}{\Delta} + L \quad \forall\beta \in \mathcal{B}, \tag{B.5}$$

and

$$\phi(s)c^R(r,y) \geq \eta p U(y)\epsilon - C_H r, \tag{B.6}$$

where L is the approximation of the differential generator as defined in (A.6).

The problem of solving for the approximating coefficients $c^R(r,y)$ $\forall y > 0$ for a given r may be represented as an extended vertical linear complementarity problem (EVLCP) due to the fact that one of the conditions in (B.5)-(B.6) must hold with equality. The EVLCP will be of the form

$$\min\left[B_1^R c^R(r,y) + b_1^R, B_2^R c^R(r,y) + b_2^R, B_3^R c^R(r,y) + b_3^R\right] = 0. \tag{B.7}$$

Given a vector of q nodal points $\hat{s}$ the EVLCP is defined by

$$B_l^R = \begin{cases} \Phi & \text{if } l = 1 \\ \rho\Phi - \hat{L} & \text{if } l = 2 \\ \left[\rho + \frac{1}{\Delta}\right]\Phi - \hat{L} & \text{if } l = 3 \end{cases},$$

and

$$b_l^D = \begin{cases} -\eta\hat{p}\pi_y\epsilon + C_H r & \text{if l=1} \\ C_N & \text{if } l = 2 \\ \frac{-1}{\Delta}\Phi c^R(r-\Delta) + C_H + C_R & \text{if } l = 3 \end{cases},$$

where Φ is the set of basis functions evaluated at the nodal points $\hat{s}$ and $\hat{L}$ is the differential generator approximation evaluated at $\hat{s}$.

Since entitlements have already been obtained, $y > 0$, when remediation is complete $r = 0$ the value will be equivalent to the development value with $\pi_y\epsilon$ units of construction remaining as defined in the boundary condition (2.17) such that

$$\Phi c^R(0, y) = \Phi c^D(\pi_y \epsilon).$$

This defines the vector of approximating coefficients, $c^R(0, y)\ \ \forall y > 0$, at the end of the initial stage. Given this information we may iterate backwards in the r space solving the EVLCP in (B.7) for the next set of approximating coefficients at each Δ step using a newton-type method and a semi-smooth version of the EVLCP.

Given the approximation for the value function after the regulation process has concluded the prior expectation may be computed using (2.16)

$$E\left[V^R(s, r, \tilde{y})\right] \approx \phi(s) \sum_{j=1}^{N} Pr(\pi_j) c^R(r, j), \tag{B.8}$$

where $\tilde{y} > 0$. The value of the project before completing remediation, $r > 0$, and obtaining entitlements, $y = 0$, will satisfy the variational inequality in (2.13)-(2.15), which may be rewritten in terms of the approximation as

$$\begin{aligned} \rho\phi(s)c^R(r, y) \geq \quad & f(y, \beta) - \beta\phi(s)\tfrac{c^R(r,y) - c^R(r-\Delta,y)}{\Delta} + L \\ & + \tfrac{1}{\lambda}\phi(s)\left[\textstyle\sum_{j=1}^{N} Pr(\pi_j)c^R(r, j) - c^R(r, 0)\right] \quad \forall \beta \in \mathcal{B}, \end{aligned} \tag{B.9}$$

and

$$\phi(s)c^R(r, y) \geq \eta p U(y)\epsilon - C_H r, \tag{B.10}$$

where L is the approximation of the differential generator as defined in (A.6). Again it is possible to rewrite the system as an EVLCP of the form

$$\min\left[\tilde{B}_1^R c^R(r, y) + \tilde{b}_1^R, \tilde{B}_2^R c^R(r, y) + \tilde{b}_2^R, \tilde{B}_3^R c^R(r, y) + \tilde{b}_3^R\right] = 0, \tag{B.11}$$

where

$$\tilde{B}_l^R = \begin{cases} \Phi & \text{if } l = 1 \\ (\rho + \frac{1}{\lambda})\Phi - \hat{L} & \text{if } l = 2 \\ \left(\rho + \frac{1}{\lambda} + \frac{1}{\Delta}\right)\Phi - \hat{L} & \text{if } l = 3 \end{cases},$$

and

$$\tilde{b}_l^D = \begin{cases} -\eta\hat{p}\pi_y\epsilon + C_H r & \text{if l=1} \\ C_N - \frac{1}{\lambda}\Phi\sum_{j=1}^{N} Pr(\pi_j)c^R(r,j) & \text{if } l=2 \\ \frac{-1}{\Delta}\Phi c^R(r-\Delta) + C_H + C_R - \frac{1}{\lambda}\Phi\sum_{j=1}^{N} Pr(\pi_j)c^R(r,j) & \text{if } l=3 \end{cases}.$$

Similar to the approach for the remediation stage after the regulation process, initially the approximating coefficients at the terminal boundary for the deterministic state variable are computed. We then iterate backwards over the r space solving the EVLCP in (B.11) for the remaining approximating coefficients. The initial coefficients at the terminal boundary $r = 0$ are determined by the value defined in (2.18)-(2.19). The boundary condition defined by the variational inequality in (2.18)-(2.19) may be approximated in a similar fashion.

C Derivation of Moments for the Price Process

The two factor process for may be described as $\tilde{S}_t = (\ln(P_t), \mu_t)$ where

$$d\tilde{S}_t = \left(\begin{bmatrix} -\frac{1}{2}\sigma_p^2 \\ \delta\bar{\mu} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} \tilde{S}_t \right) dt + \begin{bmatrix} \sigma_p & 0 \\ 0 & \sigma_\mu \end{bmatrix} dW_t,$$

where $dW = (dW_1, dW_2)$. In integral form this may be represented as

$$\tilde{S}_{t+\Delta} = \tilde{S}_t + \int_0^\Delta \left(\begin{bmatrix} -\frac{1}{2}\sigma_p^2 \\ \delta\bar{\mu} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} \tilde{S}_{t+h} \right) dh + \int_0^\Delta \begin{bmatrix} \sigma_p & 0 \\ 0 & \sigma_\mu \end{bmatrix} dW_h.$$

Noting the fact that the diffusion term will have an expectation of 0 it may be seen that

$$E\left[\tilde{S}_{t+\Delta} \middle| \tilde{S}_t\right] = \tilde{S}_t + \int_0^\Delta \left(\begin{bmatrix} -\frac{1}{2}\sigma_p^2 \\ \delta\bar{\mu} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} E\left[\tilde{S}_{t+h} \middle| \tilde{S}_t\right] \right) dh,$$

which may be rewritten as

$$\frac{\partial E\left[\tilde{S}_{t+\Delta} \middle| \tilde{S}_t\right]}{\partial \Delta} = \begin{bmatrix} -\frac{1}{2}\sigma_p^2 \\ \delta\bar{\mu} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} E\left[\tilde{S}_{t+\Delta} \middle| \tilde{S}_t\right].$$

This represents an ordinary differential equation with the boundary condition $E\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = \tilde{S}_t$ for $\Delta = 0$. The solution to which is

$$E\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = \begin{bmatrix} \bar{\mu}\Delta + \frac{1}{\delta}\left(\bar{\mu}-\mu_t\right)\left(e^{-\delta\Delta}-1\right) - \frac{1}{2}\sigma_p^2\Delta + \ln(P_t) \\ e^{-\delta\Delta}\left(\mu_t - \bar{\mu}\right) + \bar{\mu} \end{bmatrix}. \tag{C.1}$$

To find the variance of the process we consider the function $y_t = E\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = Y\left(\tilde{S}_t\right)$. By Ito's Lemma

$$\begin{aligned} dy = & \left\{ \tfrac{\partial Y}{\partial t} + \tfrac{\partial Y}{\partial \tilde{S}}\left(\begin{bmatrix} -\frac{1}{2}\sigma_p^2 \\ \delta\bar{\mu} \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix}\tilde{S}_t\right) + \tfrac{1}{2}\tfrac{\partial^2 Y}{\partial \tilde{S}^2}\mathrm{vec}\left(\begin{bmatrix} \sigma_p^2 & 0 \\ 0 & \sigma_\mu^2 \end{bmatrix}\right)\right\} dt \\ & + \tfrac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} dW. \end{aligned}$$

The drift must be equal to zero due to the law of iterated expectations and therefore this reduces to

$$dy = \frac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} dW.$$

This may be rewritten in integral form such that

$$\tilde{S}_{t+\Delta} = y_t + \int_0^\Delta \frac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} dW, \tag{C.2}$$

for which the variance is

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = E\left[\left(\tilde{S}_{t+\Delta} - y_t\right)\left(\tilde{S}_{t+\Delta} - y_t\right)^T\middle|\tilde{S}_t\right].$$

Given the definition in (C.2) this variance may be rewritten as

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = E\left[\left(\int_0^\Delta \frac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} dW\right)\left(\int_0^\Delta \frac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix} dW\right)^T\middle|\tilde{S}_t\right].$$

Which reduces to the problem of a single integral such that

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = E\left[\int_0^\Delta \left(\frac{\partial Y}{\partial \tilde{S}}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix}\begin{bmatrix} 0 & 1 \\ 0 & -\delta \end{bmatrix}^T \frac{\partial Y}{\partial \tilde{S}}^T\right)\middle|\tilde{S}_t\right]. \tag{C.3}$$

From the definition of the first moment in (C.1) it may be seen that

$$\frac{\partial Y}{\partial \tilde{S}} = \begin{bmatrix} 1 & \frac{1}{\delta}\left(1-e^{-\delta\Delta}\right) \\ 0 & e^{-\delta\Delta} \end{bmatrix}.$$

Therefore the definition of the variance in (C.3) may be rewritten as

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] = E\left[\int_0^\Delta \begin{pmatrix} \sigma_p^2 + \frac{1}{\delta}\left(1-e^{-\delta\Delta}\right)^2\sigma_\mu^2 & \frac{1}{\delta}\left(1-e^{-\delta\Delta}\right)e^{-\delta\Delta}\sigma_\mu^2 \\ \frac{1}{\delta}\left(1-e^{-\delta\Delta}\right)e^{-\delta\Delta}\sigma_\mu^2 & e^{-2\delta\Delta}\sigma_\mu^2 \end{pmatrix}\middle|\tilde{S}_t\right],$$

which has the solution

$$Var\left[\tilde{S}_{t+\Delta}\middle|\tilde{S}_t\right] == \begin{bmatrix} \sigma_p^2\Delta + \frac{\sigma_\mu^2}{\delta^3}\left[e^{-\delta\Delta}\left(2-\frac{1}{2}e^{-\delta\Delta}\right)+\frac{2\delta\Delta-3}{2}\right] & \frac{\sigma_\mu^2}{2\delta^2}\left(e^{-\delta\Delta}-1\right)^2 \\ \frac{\sigma_\mu^2}{2\delta^2}\left(e^{-\delta\Delta}-1\right)^2 & \frac{\sigma_\mu^2}{2\delta}\left(1-e^{-2\delta\Delta}\right) \end{bmatrix}.$$