The Dissertation Committee for Warren Joseph Hahn Certifies that this is the approved version of the following dissertation:

A Discrete-Time Approach for Valuing Real Options with Underlying Mean-Reverting Stochastic Processes

Committee:

James S. Dyer, Supervisor

Leon S. Lasdon

Efstathios Tompaidis

Thaleia Zariphopoulou

Genaro J. Gutierrez

Courtland Huber

A Discrete-Time Approach for Valuing Real Options with Underlying Mean-Reverting Stochastic Processes

by

Warren Joseph Hahn, B.S., M.S., M.B.A.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

May, 2005

To my family: Lori, Allison, and Lindsey

Acknowledgements

I wish to thank my dissertation advisor, Professor James S. Dyer for his insights and guidance. I would also like to recognize the other members of my dissertation committee: Dr. Leon Lasdon, Dr. Stathis Tompaidis, Dr. Thaliea Zariphopoulou, Dr. Genaro Gutierrez, and Dr. Court Huber. Their comments and service were greatly appreciated. Dr. Luiz Brandao also provided many useful insights, and Dr. Jim Smith and Dr. John Butler were very helpful with computational matters.

I would like to thank my parents, Warren and Katie Hahn for their constant encouragement and for all the sacrifices they made over the years to support my education. I am also particularly indebted to my godparents, J.T. and Carolyn Smith for their support and guidance, and to my grandmother Willodene Watkins Smith for impressing upon me at a very early age the importance of educational opportunity.

Finally, I want to express my deep appreciation to my wife Lori and daughters Allison and Lindsey for their love, encouragement, patience and understanding during both the MBA and Ph.D. programs.

A Discrete-Time Approach for Valuing Real Options with Underlying Mean-Reverting Stochastic Processes

Publication No._____________

Warren Joseph Hahn, Ph.D.

The University of Texas at Austin, 2005

Supervisor: James S. Dyer

In this research the recombining binomial lattice approach for valuing real options is generalized to address a common issue in many real valuation problems, underlying stochastic processes that are mean-reverting. Binomial lattices were first introduced to approximate stochastic processes for valuation of financial options, and they provide a convenient framework for numerical analysis. Unfortunately, the standard approach to constructing binomial lattices can result in invalid probabilities of up and down moves in the lattice when a mean-reverting stochastic process is to be approximated. There have been several alternative methods introduced for modeling mean-reverting processes, including simulation-based approaches and trinomial trees, however they unfortunately complicate the numerical analysis of valuation problems. The approach developed in this research utilizes a more general binomial approximation methodology from the existing literature to model simple homoskedastic mean-reverting stochastic processes as recombining lattices. This approach is then extended to model a two-factor mean-

reverting process that allows for uncertainty in the long-term mean, and to model two correlated one-factor mean-reverting processes. These models facilitate the evaluation of real options with early-exercise characteristics, as well as multiple concurrent options.

The models developed in this research are tested by implementing the lattice in binomial decision tree format and applying to hypothetical real option examples with underlying mean-reverting commodity price. To specify the stochastic process for commodity price, different data analysis techniques such as Kalman filtering and seemingly unrelated regression are used. These different techniques are empirically tested to evaluate differences in the estimates and assess the tradeoffs in computational requirements. To validate the binomial model, results are compared to those from simulation-based methods for simple options. The convergence properties of the model and the relationship between length of time increment and accuracy of solutions obtained are also investigated. For cases where the number of discrete time periods becomes too large to be solved using common decision tree software, recursive dynamic programming algorithms are developed to generate solutions. Finally, we illustrate a real application by solving for the value of an oil and gas switching option which requires a binomial model of two correlated one-factor commodity price models.

TABLE OF CONTENTS

List of Tables

List of Figures

1. INTRODUCTION

The seminal work of Black and Scholes (1973) and Merton (1973) in the area of financial option valuation led to the application of option pricing methods in valuing real investments under uncertainty by recognizing the analogy between financial options and project decisions that can be made after some uncertainties are resolved. This approach has the advantage of including the value of managerial flexibility, which is frequently not captured by standard valuation approaches.

The options derived from managerial flexibility are commonly called "real options" to reflect their association with real assets rather than with financial assets. Despite its theoretical appeal, however, the practical use of real option valuation techniques in industry has been limited by the mathematical complexity of these techniques and the resulting lack of intuition associated with the solution process, or the restrictive assumptions required to obtain analytical solutions.

The mathematical complexity associated with option theory stems from the fact that the general problem requires a probabilistic solution to a firm's optimal investment decision policy at the present time and also at all instances in time up to the maturity of its options. To solve this problem of dynamic optimization, the evolution of uncertainty in the value of the real asset over time is first modeled as a stochastic process. Then the value of the firm's optimal policy over time is obtained as the solution to a stochastic differential equation with appropriate boundary conditions to reflect the initial conditions and terminal payoff characteristics. Recursive dynamic programming may be used to obtain closed-form mathematical solutions for certain types of stochastic processes and for specific exercise characteristics of options.

However, a complicating key assumption is that market-based information (i.e., information on prices of traded assets) can be used to specify the parameters of the stochastic process for the underlying asset. If not, as is often the case with options on real projects, there is no information on the market's view of the risk associated with the project. Hence, there is no market-based guidance for selecting the discount rate to be applied to cash flows.

To provide a transparent, computationally efficient model of the valuation problem, a discrete approximation of the underlying stochastic process can be developed. The first example of this approach was a binomial lattice model that converges weakly to a lognormal diffusion of stock prices known as a Geometric Brownian motion or GBM, developed by Cox, Ross, and Rubinstein (1979). The binomial model can be used to accurately approximate solutions from the Black-Scholes-Merton continuous-time option valuation model. Moreover, this approach can also be used to solve for the value of early-exercise American options, whereas the Black-Scholes-Merton model can only value European and infinite-horizon American options.

However, the assumption of a lognormal geometric Brownian diffusion as a model of the underlying stochastic process may not be valid for many real option valuation problems, such as projects with cash flows that depend on mean-reverting commodity prices. The effect of modeling a stochastic process that is mean-reverting with a lognormal geometric Brownian diffusion model can be a significant overestimation of uncertainty in the resultant cash flows from a project, which can result in overstated option values. Figure 1.1 shows a comparison of GBM and mean-reverting diffusions with the same standard deviation of returns.

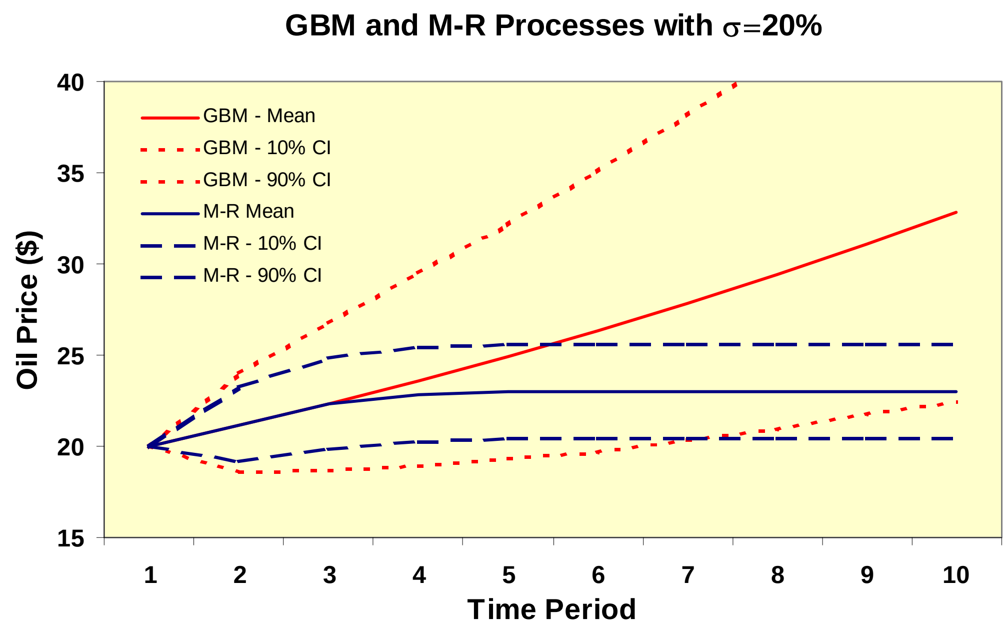

Figure 1.1 – Comparison of GBM and Mean-Reverting Diffusions

Discrete-time modeling of mean-reverting stochastic processes has proven problematic, however. Methods employing Monte Carlo simulation and discrete trinomial trees have been the two primary proposed approaches. Unfortunately these methods are computationally intensive and difficult to implement for the more complex problems encountered in real options.

In this research the method for constructing recombining binomial lattices developed by Nelson and Ramaswamy (1990) is extended to develop binomial models for homoskedastic mean-reverting stochastic processes, including the two-factor model of

Schwartz and Smith (2000). A goal of this research is to demonstrate that these models can be used for real option problems, such as those used to model commodity price. This approach can be implemented in binary decision trees with off-the-shelf decision tree software.

This dissertation is organized as follows: Section 2 contains a review of the relevant literature for this topic. In Section 3, the Nelson and Ramaswamy (1990) approach to constructing computationally simple binomial lattices is reviewed for the case of a one-factor mean-reverting process and then extended to model both two-factor mean-reverting diffusions and two correlated one-factor diffusions. Section 4 details how a two-factor model can be implemented in decision tree and lattice formats, and investigates the model's convergence properties numerically for the two-factor diffusion model of Schwartz and Smith (2000) up to the computational limits of decision tree software, and from that point forward with a coded lattice algorithm. In Section 5, the different methods for determining the parameters for the mean-reverting processes to be modeled are presented, and the results from application of each approach to an extensive futures data set are discussed. In Section 6, the approach developed in this research is applied to a real example of a switching option in an oil and gas setting which requires a binomial model of two correlated one-factor models. Finally, in Section 7, conclusions from this work and further research issues regarding model formulation and application to real problems are discussed.

2. LITERATURE REVIEW

Discounted cash flow methods (DCF) are commonly used in practice for the valuation of projects and for decision-making regarding investments in real assets. Under this approach, the value of a project is determined by discounting the future expected cash flows at a discount rate that reflects the riskiness of the project. In practice, most projects are valued using the weighted average cost of capital for the firm, or WACC, as the discount rate. This assumes the project's risks are essentially equal to the risks associated with the firm as a whole, which may not be appropriate for many investment projects. The DCF approach also assumes that once the firm commits to a project, the project's outcome will be unaffected by future decisions, thereby ignoring any managerial flexibility the project may have. Option pricing approaches can be used to address the shortcomings of the traditional DCF approach and provide an integrated approach to risk and its effect on value.

2.1 OPTION PRICING TECHNIQUES

Option pricing approaches are founded in the work of Black and Scholes (1973) and Merton (1973) in the area of financial option valuation. Traditional option pricing methods are based on the concept of no-arbitrage pricing and therefore require that markets be complete. In complete markets, there are a sufficient number of traded assets to allow the creation of a portfolio of securities whose payoffs exactly replicate the payoffs of the asset in all states of nature and in all future periods. Rubinstein (1976) and Brennan (1979) also showed that if the return on an asset is lognormal, under the assumptions of aggregation and constant proportional risk aversion the Black-Scholes formula holds even without the ability to construct a riskless hedge. This is important in cases where the basic exogenous variables are cash flows from assets, as in the case of many corporate finance problems, rather than the value of a traded asset.

A shortcoming of the Black-Scholes-Merton model and most continuous-time closed-form solutions for option value is that only options exercised at maturity, so-called European options, can be valued. Geske and Johnson (1984) developed an analytical expression for the value of American put options and proved that it holds in the limit, however it cannot be directly evaluated, as the solution to the partial differential equation is subject to boundary conditions at an infinite number of discrete points. Consequently, they proposed picking discrete evaluation points and extrapolating results to obtain the value estimate. Ramaswamy and Sundaresan (1985) note in their work on interest rate derivatives that, in some cases, the incremental values due to early exercise of American options are small, and therefore the European price serves as a useful approximation. Unfortunately, there are no general rules for when this approximation might be adequate, and there are certainly applications in which such an approach would be unsatisfactory.

To value American options and other types of options that can be exercised before maturity, numerical techniques are typically used. The binomial approximation of Sharpe (1978) and Cox, Ross, and Rubinstein (1979) and finite difference methods are the two primary techniques that have been introduced for this purpose.

Binomial models are accurate, remarkably robust, and intuitively appealing tools for valuing financial and real options. A well-known example of specifying parameters for a recombining binomial lattice is the model of the stochastic differential equation

$\frac{dS}{S} = \mu dt + \sigma dz$, which represents a GBM model of the diffusion of asset price over time. Using an important result from stochastic calculus, Ito's Lemma, we can write the corresponding transformed process for the log of asset price as

$$d \ln S = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dz ,$$

where S is the asset price (eg., stock price), μ is the growth rate (drift), σ is the standard deviation of returns (volatility), and dz is a Wiener process (random increment with mean zero and variance of dt).

By requiring that the first and second moments of a binomial distribution match those of the continuous diffusion, the up and down movements at each step in a lattice are calculated to be $u = e^{\sigma \sqrt{\Delta t}}$ and $d = e^{-\sigma \sqrt{\Delta t}}$, respectively, and the probabilities of the up and down movements are $p = \frac{1 + \mu - d}{u - d}$ and $1 - p$, respectively. The asset price in period i and state j is $S_{i,j} = S_0 u^{i-j} d^{j}$. This model converges weakly to the above GBM as the time increment Δt approaches zero. As a result it has been very popular in valuing financial options and many types of real options where a GBM is a reasonable representation of the diffusion of the underlying asset value. Furthermore, binomial models can be used to implement either the riskless hedge approach to option valuation, using the Cox, Ross, Rubinstein model, or the preference-based approach of Rubinstein (1976) and Brennan (1979), as outlined by Stapleton and Subrahmanyam (1984).

Binomial approximations have also been developed for two-factor diffusions (Boyle, 1988) and have been used extensively in modeling interest rate dynamics. Boyle, Evnine, and Gibbs (1989) illustrate the use of a generalized lattice framework for multivariate contingent claims by computing option values and checking against values from closed-form solutions. Madan, Milne, and Shefrin (1989), He (1990), and Ho, Stapleton, and Subrahmanyam (1995) have also demonstrated that the convergence of these models duplicates that of the univariate case for GBM diffusions. Amin (1991) extends the discrete binomial approximation for both univariate and multivariate cases to allow for time-varying volatility functions. The time-varying volatility is accommodated in this approach by introducing a time-dependent step size that offsets changing volatility.

The typical approach for solving for option value using a binomial lattice is to find the replicating portfolio at each node, working backwards through the lattice. Unfortunately, this process can be cumbersome and non-intuitive, especially for more complex applications to real assets, which can involve several simultaneous and compound options, or involve path-dependant options.

Finite difference methods were first introduced for option valuation by Schwartz (1977) and later extended to value exercise options with jump diffusion stochastic processes (Brennan and Schwartz, 1978). While finite difference methods have the advantage of more flexibility in modeling underlying stochastic processes, these methods can be computationally intensive. Geske and Shastri (1985) provide a comparison of alternative option valuation methods, including a binomial model and several different finite difference models, based on both accuracy and computational time. Their results demonstrate the binomial models run in a fraction of the time required for most finite difference models and are generally more stable, although finite difference models were more accurate in some cases. They also note that binomial models are pedagogically superior.

2.2 MEAN REVERTING PROCESSES

The assumption of a lognormal geometric Brownian diffusion as a model of the underlying stochastic process may not be valid for many problems. This is a key issue, as pointed out by Cox and Ross (1976), who note the importance of the specification of underlying stochastic process in valuation of options by reviewing the assumptions employed in the Black-Scholes model and evaluate alternative forms of processes. Their study of so-called "single-stage" jump processes was a necessary prelude to the subsequent development of their binomial model.

Many valuation problems have underlying stochastic processes that are mean-reverting, such as projects with cash flows that depend on mean-reverting commodity prices. Most empirical studies of historical commodity data have found that mean-reverting models accurately capture the evolution of prices (e.g., Schwartz, 1997). Bessembinder, et al. (1995) find that a forward-looking analysis of the commodities futures data implies mean reversion as well. There are a few empirical studies of commodity data that do not support the mean-reverting hypothesis, but they are either for special cases or are inconclusive. Hjalmarsson (2003) finds that electricity option prices based on the GBM assumption are more accurate in matching non-parametric estimates than are prices calculated from an Ornstein-Uhlenbeck mean-reverting process, however it is likely that this analysis may be affected by the lack of an efficient electricity options trading market, as well as by the fact that electricity is essentially a non-storable commodity.

Bhattacharya (1978) demonstrates that mean-reverting cash flows are in general likely to be more realistic for many investment projects in a competitive economy, since the expectation is that cash flows from a particular project will revert to levels that make firms indifferent about new investments of the same type. Metcalf and Hassett (1995) study investment under the assumptions of lognormality and mean-reversion and find offsetting consequences. Under the lognormal assumption investments derive value from the option effect and the possibility of higher future payoffs, whereas under mean-reversion it is the reduced risk that encourages investment. Lo and Wang (1995) show that drift indirectly affects options prices, and thus predictability in returns and mean-reversion will affect option prices. They demonstrate this by comparing option prices for a hypothetical stock under lognormal and mean-reverting Ornstein-Uhlenbeck processes assumptions. In order to do this, they set the distribution of the underlying process and find its implications for the risk-neutral process. This is the reverse approach of Gundy (1991) who takes the risk-neutral distribution reflected in derivative prices and infers the properties of the true process.

As noted by Schwartz (1998), Laughton and Jacoby (1993), and others, if commodity prices are indeed mean-reverting, then a lognormal geometric Brownian diffusion model can significantly overestimate uncertainty in the resultant cash flows from a project, and result in overstated option values.

2.3 NUMERICAL TECHNIQUES

In cases where the underlying stochastic process should be modeled as mean-reverting, rather than as a GBM, the problem can be solved in one of two ways: 1) use a Monte Carlo simulation method, thereby eliminating the need to build a tree to represent the stochastic process or 2) use a different type of tree-building procedure or finite difference approach. Monte Carlo methods are straightforward to apply for European options, and can also be used to value American options in some cases. Longstaff and Schwartz (2001) proposed a method employing Monte Carlo simulation which can accommodate general types of stochastic processes, and can also be used to value early-exercise options. This method uses ex-post regression of cash flows on state values at each step to estimate the value function used to determine the optimal stopping rule, and hence option value. However, a significant drawback of this approach is that it is computationally intensive, non-intuitive, and limited to a small number of relatively simple types of project options.

Other researchers have developed discrete tri- and multi-nomial trees for valuing options in a similar manner to the binomial approach, but with the ability to model more general types of stochastic processes, due to the additional degrees of freedom. Hull and White (1990a) introduce the approach whereby the initial term structure of futures prices is matched by including an adjustment term, $\theta(t)$ in the diffusion equation. In this paper they use this approach to extend two different mean-reverting interest rate models, Vasicek (1977) and the Cox, Ingersoll, and Ross (1985), so that the initial term structure is exactly matched. They integrate this with their work on valuing derivatives using the explicit finite difference method (Hull and White, 1990b) to produce a procedure for valuing derivative securities that have underlying mean-reverting interest rate processes (Hull and White, 1993a). In this paper they assert that recombining binomial models cannot be used in general to model these types of processes, but acknowledge in a footnote that the approach of Nelson and Ramaswamy (1990) can be used if the expected drift and variance at each step are required to be correct only in the limit. They do not comment on whether desired levels of accuracy might be obtained within a reasonable number of steps. In this approach, the values of the adjustment θ and drift μ are assumed to be constant in between the increments, and the length of the increments is set by the frequency of futures maturities. This approach is later improved to provide faster tree construction, more accurate pricing, and better convergence by changing the geometry of the trinomial tree so that the central node at each step corresponds to the expected value (Hull and White, 1993b).

Hull and White (1994b) also show that their approach can be extended to model two-factor processes or two correlated one-factor processes, and illustrate for the example of interest rate derivatives from two countries. The bivariate Hull and White approach entails calibration of two separate trees and construction of a combined tree with nine branches emanating from each node. To adjust for correlation, nine factors must be calculated for each node to adjust the branching probabilities, and further, the calculation of the adjustments depends on whether the correlation is positive, zero, or negative. Probabilities can still be negative at some nodes under this procedure, so the adjustment factor is set to the maximum value for which probabilities are non-negative. Hull and White acknowledge that this introduces some bias in correlation, but claim that this bias disappears as the time increment goes to zero. Some of the other issues with the use of univariate approach, including variable time steps, cash flows that occur between nodes, and path dependency, and more detail on use of the method are provided in a subsequent paper (Hull and White, 1996). The approach suggested for interim cash flows is either to discount to the nearest node or to increase the number of time steps so that nodes occur at the same frequency as cash flows. They do not describe how to calibrate such a tree at points where no futures data exists.

The difficulties with implementing the Hull-White multivariate method and issues with performance are discussed in both Muck and Rudolf (2002) and Staley and Wicentowich (2003). In the latter of these, the authors propose as an alternative a tree in which the probabilities and node spacing are set to match the volatility structure, and the drift term is allowed to be miss-specified. The drift errors, or difference between the true process drift and the miss-specified drift, are stored for each node, and the tree is then adjusted during backward induction to calculate option values. Derivative values must then be calculated by cubic spline interpolation to compensate for underlying drift errors. These authors acknowledge the computational burden with this approach, especially with the Hull-White two factor diffusion. These authors also cite the Nelson and Ramaswamy (1990) approach as an alternative, but note that it has not been extended to the multivariate case.

Other two-factor mean-reverting processes include Schwartz (1997), Schwartz and Smith (2000), Gibson and Schwartz (1990), and Ribeiro and Hodges (2004). These composite diffusions generally include a second factor to explicitly model uncertainty in the short-term deviations from the long-term mean, as well as in the long-term mean itself. None of these diffusions have yet been approximated with a discrete recombining lattice or tree.

Given the difficulty in implementing trinomial trees, there is a need for a modeling procedure similar to the binomial approach to a GBM that exhibits convergence in distribution for other distributions, and is yet computationally simple and robust in terms of allowable payoff specifications.

Nelson and Ramaswamy (1990) propose a modeling procedure that exhibits convergence in distribution under very general conditions. The binomial sequence of Cox, Ross, and Rubinstein is in fact a special case of this procedure. For diffusions with constant variance, this approach entails fixing the up and down moves in the tree and calculating probabilities at each node, conditioned on the state, to reflect the local drift. In any cases where nodes have invalid probabilities, the probabilities are censored so that negative probabilities are set to zero and probabilities greater than one are set to one. Nelson and Ramaswamy show that as the time step is reduced the drift and variance of this approximation converge to those of the continuous diffusion. This paper is preceded by Nelson's (1990) investigation of the use of discrete time ARCH stochastic difference equation systems to approximate continuous diffusions, in which conditions for a finite dimensional discrete time Markov process to converge to an Ito process are presented.

This method can also be applied, with additional calibration steps, to the development of recombining lattices for heteroskedastic stochastic processes. This additional step is required to calibrate the up and down moves to reflect the local variance, and basically entails a transformation of the process to remove the heteroskedasticity. Subsequent work with this method has centered on developing models for valuing interest rate derivatives. Peterson, Stapleton, and Subrahmanyam (1999) provide an example of this line of research. This model also appears in work to develop a discrete model to value American options when the underlying uncertainty follows a jump diffusion process (Amin, 1993), and in work to develop a discrete-time model to price currency exchange rate derivatives with stochastic volatility (Amin and Bodurtha, 1995).

Other proposed methods for constructing recombining lattices for general types of stochastic processes, such as the variable jump approach proposed by Calistrate, Paulhus, and Sick (1999) and a lattice based on an inhomogeneous geometric Brownian motion (Robel, 2001) also appeal to the Nelson and Ramaswamy approach to calculating probabilities of up and down moves in the lattice.

2.4 Real Options and Decision Analysis

Building on the success of the Black-Scholes-Merton and Cox-Ross-Rubinstein approaches in the area of financial option valuation, option pricing methods were soon applied to the valuation of real investments under uncertainty to address the shortcomings of the traditional DCF approach. The fundamental premise, as pointed out by Rubinstein (1994) is that asset prices in efficient markets contain valuable information that can be used in making economic decisions. Some of the first examples were Tourinho (1979), who used the concept of an option to evaluate a non-renewable natural resources reserve under price uncertainty; Brenann and Schwartz (1985), who analyzed the optimal operational policy of a copper mine; and McDonald and Siegel (1986), who determined the optimal timing for investing in a project with irreversible investments with uncertain cost and benefits. Dixit and Pindyck (1994) and Trigeorgis (1996) were among the first authors to synthesize several of these ideas. Most of the early applications were either attempts to adapt continuous-time analytical solutions or lattice-based approaches similar in spirit to the Cox-Ross-Rubinstein approach using replicating portfolios. However, most projects involving real assets do not have a replicating portfolio of securities, so markets are not complete. In this case, Dixit and Pindyck (1994) propose the use of dynamic programming using a subjectively defined discount rate, but the result does not provide a true market value for the project and its options.

The application of decision analysis to real option valuation problems seems natural because decision trees are commonly used to model project flexibility, but there has only been limited work in this area (Howard (1996)). Nau and McCardle (1991) and Smith and Nau (1995) study the relationship between option pricing theory and decision analysis and demonstrate that the two approaches yield the same results when applied correctly. Smith and Nau propose a method which integrates the two approaches by distinguishing between market risks, which can be hedged by trading securities and valued using option pricing theory, and private uncertainties which are project-specific risks and can be valued using decision analysis techniques. Smith and McCardle (1998, 1999) illustrate how this approach can be applied in the context of oil and gas projects, and provide a discussion of lessons learned from applications to some case studies.

To transition from a binomial lattice to a probability tree, a tree is constructed with binary chance branches that have the unique feature that the outcome resulting from moving up and then down in value is the same as the outcome from moving down and then up. For example, Figure 2.1 shows a binomial lattice, along with a binary tree that has the relationship between up and down movements at each node specified as $up = \frac{1}{down}$. From this figure, it is evident that there will be the same number of different outcomes in any period, although some of the outcomes will be recurring in the binomial tree.

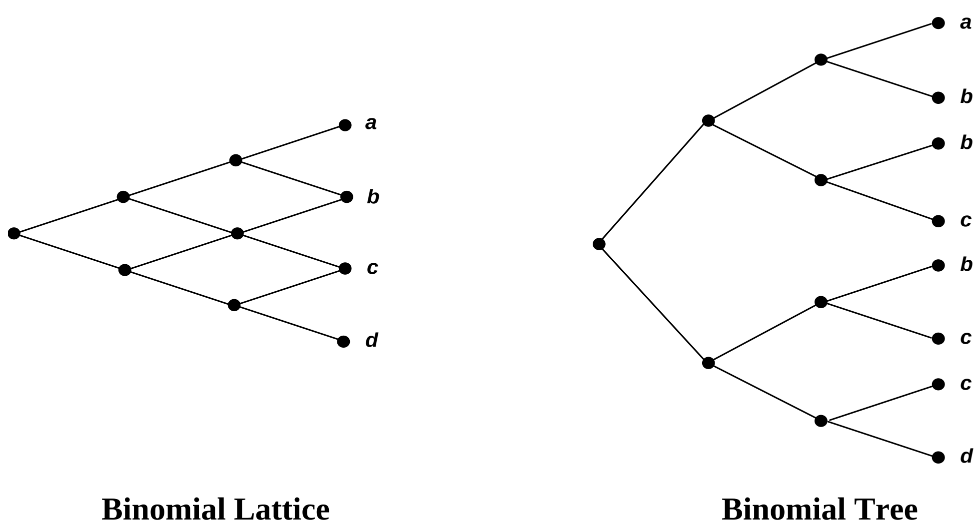

Figure 2.1 – Comparison of Recombining Binomial Lattice and Binomial Tree

To value options in this format, decision nodes are added at each point in the tree where exercise decisions exist, with corresponding payoffs entered at each terminal node (Brandao and Dyer, 2004; Brandao, Dyer, and Hahn, 2004). The binomial model has the important property of recombination, that is, branches of the binomial lattice reconnect at each step. This is an important issue from a computational perspective, because there are $N+1$ nodes at any stage N, whereas there are 2^N nodes at the same stage for a binary tree. Therefore, problems with large values of N may require algorithms that are coded to take advantage of efficiencies provided by recombining lattices.

3. GENERAL METHOD OF DEVELOPING RECOMBINING LATTICES

In constructing the Nelson and Ramaswamy (1990) model, the problem is to find a binomial sequence that converges to a stochastic differential equation (SDE) of the general form:

$$dY_t = \mu(Y,t)dt + \sigma(Y,t)dz$$

where $\mu(Y,t)$ and $\sigma(Y,t)$ are continuous instantaneous drift and standard deviation functions, and dz is a standard Brownian increment. To solve this problem, Nelson and Ramaswamy propose a simple binomial sequence of n periods of length Δt, where $\Delta t = \frac{T}{n}$, and

$$Y_t^+ \equiv Y + \sqrt{\Delta t}\sigma(Y,t) \qquad \text{(up move)}$$

$$Y_t^- \equiv Y - \sqrt{\Delta t}\sigma(Y,t) \qquad \text{(down move)}$$

$$q_t \equiv \frac{1}{2} + \sqrt{\Delta t}\,\frac{\mu(Y,t)}{2\sigma(Y,t)} \qquad \text{(probability of up move)}$$

$$1 - q_t \qquad \text{(probability of down move)}$$

The conditions under which this sequence converges to the above SDE are 1) that the SDE is well-behaved (i.e., that $Y_t = Y_0 + \int_0^t \mu(Y_s,s)ds + \int_0^t \sigma(Y_s,s)dz_s$ exists on $0 < t < \infty$) and 2) that the jump sizes, local drift, and local variance converge in distribution (i.e., $\left|Y_t^{\pm}(Y,t)-Y\right|$, $\left|\mu_t(Y,t)-\mu(Y,t)\right|$, and $\left|\sigma_t^2(Y,t)-\sigma^2(Y,t)\right| \to 0$ as $\Delta t \to 0$).

3.1 ONE-FACTOR MEAN-REVERTING MODELS

This approach can be applied to a mean-reverting process to facilitate the evaluation of real options on commodity price-contingent projects. First consider a simple one-factor mean-reverting process, the Ornstein-Uhlenbeck process, which is given by:

$$dY_t = \kappa(\overline{Y} - Y_t)dt + \sigma dz_t,$$

where Y_t is the log of commodity price, κ is a mean reversion coefficient, $\overline{Y}$ is the log of long-term mean price, σ is the process volatility, and dz is a Wiener process (random increment with mean zero and variance of dt). We use the log since it is commonly assumed that commodity prices are lognormally distributed.

Substituting $\kappa(\overline{Y} - Y_t)$ for $\mu(Y,t)$ and σ for $\sigma(Y,t)$ in the above binomial sequence yields the following parameterization for the binomial model:

$$Y_t^+ \equiv Y + \sqrt{\Delta t}\sigma \qquad \text{(up move)}$$

$$Y_t^- \equiv Y - \sqrt{\Delta t}\sigma \qquad \text{(down move)}$$

$$q_t \equiv \begin{cases} \dfrac{1}{2} + \sqrt{\Delta t}\dfrac{\kappa(\overline{Y} - Y_t)}{2\sigma} & \text{if } 0 \le \dfrac{1}{2} + \sqrt{\Delta t}\dfrac{\kappa(\overline{Y} - Y_t)}{2\sigma} \le 1 \\ 0 & \text{if } \dfrac{1}{2} + \sqrt{\Delta t}\dfrac{\kappa(\overline{Y} - Y_t)}{2\sigma} \le 0 \\ 1 & \text{if } 1 \le \dfrac{1}{2} + \sqrt{\Delta t}\dfrac{\kappa(\overline{Y} - Y_t)}{2\sigma} \end{cases} \qquad \text{(probability of up move)}$$

$1 - q_t$ (probability of down move)

This specification shows the conditioning of probabilities on the deviation of the mean at each node, and the necessary censorship to values between 0 and 1. The above formula can be rewritten in one statement as:

$$q_t = \max\left(0, \min\left(1, \left(\frac{1}{2} + \sqrt{\Delta t}\,\frac{\kappa(\overline{Y} - Y_t)}{2\sigma}\right)\right)\right)$$

Thus, all of the information for modeling a one-factor mean-reverting process as a discrete-time binomial lattice or binomial tree is given.

3.2 Two-Factor Mean-Reverting Models

Although the one-factor model can be used to capture mean reversion in a parameter such as a commodity price, it assumes there is no uncertainty in the long-term mean. Gibson and Schwartz (1990), Schwartz (1997), Schwartz and Smith (2000) and others have introduced composite diffusions that include a second factor to explicitly model uncertainty in the short-term deviations from the long-term mean, as well as in the long-term mean itself. A goal of this research was to develop a discrete binomial representation of a two-factor model in a similar manner to the binomial approximations of two correlated GBM diffusions introduced by Boyle (1988). Hull and White (1994b) also show that their approach can be extended to model two-factor processes or two correlated one-factor processes, but there are several computational difficulties with this approach as was discussed in Chapter 2. We show that the general approach for tree construction discussed thus far can be extended to this two-factor model so that one or both of the factors can follow a mean-reverting diffusion.

The Schwartz and Smith diffusion (2000) is the best candidate for discrete modeling because the two factors are split apart, rather than having one factor nested in the process for the other. In this diffusion, the logarithm of the price at any point is decomposed into two factors; a long-term equilibrium price, ξ_t, and a deviation from the equilibrium price, χ_t. The long-term equilibrium price is specified to follow a GBM, while the short-term deviation follows a simple one-factor Ornstein-Uhlenbeck process and eventually reverts to zero. The price is therefore given as:

$$Y_t = e^{\chi_t + \xi_t},$$

where the two processes are:

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \qquad \text{(long-term mean price)}$$

$$d\chi_t = \kappa(0 - \chi_t)dt + \sigma_\chi dz_\chi \qquad \text{(deviation from long-term mean price).}$$

The relationship between the increments of the two processes is given by:

$$dz_\xi dz_\chi = \rho_{\xi\chi} dt$$

Thus, the correlation $\rho_{\xi\chi}$ describes the degree to which the increments move in the same $(0 < \rho_{\xi\chi} \le 1)$ or opposite $(-1 \le \rho_{\xi\chi} < 0)$ directions.

A two-dimensional binomial approximation can be developed for this process, which results in a four-branch chance node for each discrete period, as shown in Figure 3.1.

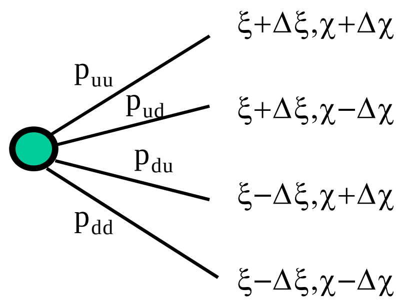

Figure 3.1 – Four-branch Chance Node for Two-factor Process

The probabilities for the joint lognormal-Ornstein-Unlenbeck process can be derived by first denoting the drift of the respective processes as

$$\nu_\xi = \mu_\xi - \frac{\sigma_\xi^2}{2} \qquad \text{(GBM for long-term mean, } \xi_t\text{)}$$

$$\nu_\chi = \kappa(0 - \chi_t) \qquad \text{(mean-reverting process for deviation, } \chi_t\text{)}$$

and selecting equal up and down jump sizes for each process:

$$\Delta_\xi = \sigma_\xi \sqrt{\Delta t} \qquad \text{(for long-term mean, } \xi_t\text{)}$$

$$\Delta_\chi = \sigma_\chi \sqrt{\Delta t} \qquad \text{(for deviation, } \chi_t\text{)}$$

Then by using the same basic method employed by Boyle (1988) for a dual lognormal approximation, we solve for the probabilities of the four possible combined outcomes by next matching the mean and variance of a two-variable binomial process. This results in the following four equations:

$$E\left[\Delta_\xi\right] = (p_{uu} + p_{ud})\Delta_\xi - (p_{du} + p_{dd})\Delta_\xi = \nu_\xi \Delta t$$

$$E\left[\Delta_\xi^2\right] = (p_{uu} + p_{ud})\Delta_\xi^2 - (p_{du} + p_{dd})\Delta_\xi^2 = \sigma_\xi^2 \Delta t$$

$$E\left[\Delta_\chi\right] = \left(p_{uu} + p_{du}\right)\Delta_\chi - \left(p_{ud} + p_{dd}\right)\Delta_\chi = \nu_\chi \Delta t$$

$$E\left[\Delta_\chi^2\right] = \left(p_{uu} + p_{du}\right)\Delta_\chi^2 - \left(p_{ud} + p_{dd}\right)\Delta_\chi^2 = \sigma_\chi^2 \Delta t$$

Adding an additional equation for the correlation;

$$E\left[\Delta_\xi \Delta_\chi\right] = \left(p_{uu} - p_{ud} - p_{du} - p_{dd}\right)\Delta_\xi \Delta_\chi = \rho \sigma_\xi \sigma_\chi \Delta t$$

and also requiring that the probabilities sum to unity;

$$p_{uu} + p_{ud} + p_{du} + p_{dd} = 1$$

yields six equations and six unknowns ($p_{uu}, p_{ud}, p_{du}, p_{dd}, \Delta_\xi, \Delta_\chi$). Solving gives the following joint probabilities.

$$p_{uu} = \frac{\Delta_\xi \Delta_\chi + \Delta_\chi \nu_\xi \Delta t + \Delta_\xi \nu_\chi \Delta t + \rho \sigma_\xi \sigma_\chi \Delta t}{4 \Delta_\xi \Delta_\chi}$$

$$p_{ud} = \frac{\Delta_\xi \Delta_\chi + \Delta_\chi \nu_\xi \Delta t - \Delta_\xi \nu_\chi \Delta t - \rho \sigma_\xi \sigma_\chi \Delta t}{4 \Delta_\xi \Delta_\chi}$$

$$p_{du} = \frac{\Delta_\xi \Delta_\chi - \Delta_\chi \nu_\xi \Delta t + \Delta_\xi \nu_\chi \Delta t - \rho \sigma_\xi \sigma_\chi \Delta t}{4 \Delta_\xi \Delta_\chi}$$

$$p_{dd} = \frac{\Delta_\xi \Delta_\chi - \Delta_\chi \nu_\xi \Delta t - \Delta_\xi \nu_\chi \Delta t + \rho \sigma_\xi \sigma_\chi \Delta t}{4 \Delta_\xi \Delta_\chi}$$

This model is the synthesis of two processes that can be approximated with recombining lattices, and therefore it is also recombining. However, as was the case with the Ornstein-Uhlenbeck approximation shown earlier, it may be necessary to censor probabilities when the degree of mean reversion required from a particular state results in probabilities greater than one (upward force of reversion) or less than zero (downward force of reversion) in a binomial node. Unfortunately, for a four-branch node for a joint process, it is not possible to directly censor the probabilities as previously described. Therefore, the approximation must be revised to address the limitations of the mean-reverting process approximation, while still retaining the capability of modeling two correlated processes.

The solution to this problem is a straightforward application of Bayes' Rule, which describes the relationship between joint, marginal, and conditional distributions. If the conditional probabilities for the binomial diffusion of χ can be derived, then the joint process can be expressed as the product of the marginal binomial process for ξ and the conditional binomial process for χ:

$$p(\xi_t \cap \chi_t) = p(\chi_t | \xi_t) p(\xi_t) \qquad \text{(Bayes' Rule)}$$

Since the joint probabilities have already been derived, the conditional probabilities for χ can be obtained by dividing by the marginal probabilities for ξ,

$$p_u = \frac{1}{2} + \frac{1}{2}\frac{\nu_\xi \Delta t}{\Delta_\xi}$$

$$p_d = \frac{1}{2} - \frac{1}{2}\frac{\nu_\xi \Delta t}{\Delta_\xi},$$

which yields:

$$p_{u|u} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi + \rho \sigma_\xi \sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|u} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi - \rho \sigma_\xi \sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{u|d} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\rho \sigma_\xi \sigma_\chi - \Delta_\chi \nu_\xi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|d} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) - \Delta t\left(\Delta_\chi \nu_\xi + \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}.$$

This formulation can be represented in decision tree format as a two-node sequence. As shown in the following schematic, the first node is a binomial node for the GBM process for the long-term mean ξ, followed by a binomial node for the conditional process for the short-term deviation χ.

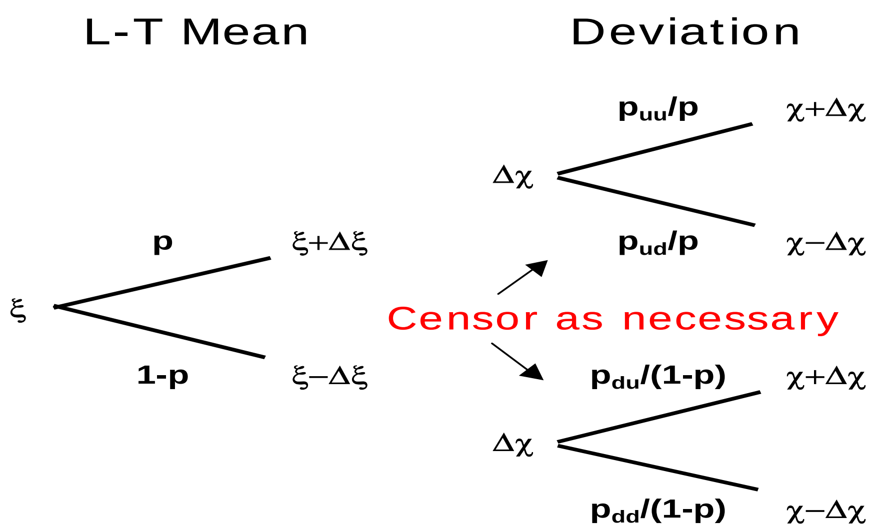

Figure 3.2 – Splitting the Four-branch node into Marginal and Conditional Steps

To check to see that the binomial approximation converges to the general SDE:

$$dY_t = \mu(Y,t)dt + \sigma(Y,t)dz,$$

the conditions are:

1) The functions $\mu(Y,t)$ and $\sigma(Y,t)$ are continuous and $\sigma(Y,t)$ is non-negative.

2) A solution of $Y_t = Y_0 + \int_0^t \mu(Y_s,s)ds + \int_0^t \sigma(Y_s,s)dz_s$ exists on $0 < t < \infty$ (*this and condition 1 ensure that the limiting SDE is well-behaved*). When this condition is satisfied, the process $\{Y_t\}_{0\le t<T}$ is characterized by:

1) the starting point Y_0

2) the continuity of Y_t

3) the drift $\mu(Y,s)$, and

4) the diffusion, $\sigma^2(Y,s)$

Given this characterization, convergence of the discrete process $\{Y_t\}$ is proved by showing the following: 1) the starting point for each increment, $Y_{0,t} \to Y_0$, 2) Jump sizes of Y_0 become small at a sufficiently rapid rate, 3) $\mu_t(Y,s) \to \mu(Y,s)$, and 4) $\sigma_t^2(Y,s) \to \sigma^2(Y,s)$. Stroock and Varadhan (1979) contains a detailed discussion of convergence requirements for discrete diffusion processes.

Since the two-factor process has been decomposed into an ABM process for ξ and a conditional arithmetic Ornstein-Uhlenbeck process for χ, convergence must be shown for both approximations. The proof of convergence for the binomial approximation of an ABM was shown by Cox, Ross, and Rubinstein (1979), and is therefore not discussed here.

For the conditional Ornstein-Uhlenbeck process, i) is satisfied, since the starting value does not change for our discrete process. For ii) to be satisfied both jumps need to converge as $\Delta t \to 0$, or:

$$\lim_{\Delta t \to 0} \sup_{\substack{|y| \le \delta \\ 0 \le t < T}} \left| Y_t^+(Y,t) - y \right| = 0 \qquad \delta > 0$$

$$\lim_{\Delta t \to 0} \sup_{\substack{|y| \le \delta \\ 0 \le t < T}} \left| Y_t^-(Y,t) - y \right| = 0 \qquad \delta > 0$$

In our particular case, $\lim_{\Delta t \to 0} \chi_t = \lim_{\Delta t \to 0} \left(\chi_{t-1} \pm \left(\sigma_\chi \sqrt{\Delta t} \right) \right) = \chi_{t-1}$, so are both satisfied. For iii) and iv), we need:

$$\lim_{\Delta t \to 0} \sup_{\substack{|y| \le \delta \\ 0 \le t < T}} \left| \mu_t(Y,t) - \mu(Y,t) \right| = 0 \qquad \delta > 0 \quad \text{and}$$

$$\lim_{\Delta t \to 0} \sup_{\substack{|y| \le \delta \\ 0 \le t < T}} \left| \sigma_t^2(Y,t) - \sigma^2(Y,t) \right| = 0 \qquad \delta > 0 \text{, respectively.}$$

The drift for the conditional process is $\kappa(\overline{Y} - Y_t) + \rho\sigma$, or in our particular case, $\kappa(0 - \chi_t) + \rho_{\xi\chi}\sigma_\chi$. However, recall that it may be necessary to censor probabilities when the degree of mean reversion required from a particular state results in probabilities greater than one (upward force of reversion) or less than zero (downward force of reversion) in a binomial node. This has an effect on the drift of the process; thus we have:

$$\mu_t(\chi,t) \equiv \begin{cases} \kappa\chi_t + \rho\sigma & if \quad 0 < p_t < 1 \\ \sigma/\sqrt{\Delta t} & if \quad p_t = 1 \\ -\sigma/\sqrt{\Delta t} & if \quad p_t = 0 \end{cases}$$

To evaluate the convergence of $\mu_t(Y,t)$, we test the limiting behavior of each of the conditional probabilities derived earlier:

$$p_{u|u} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi + \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|u} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\Delta_\chi \nu_\xi - \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{u|d} = \frac{\Delta_\xi\left(\Delta_\chi - \Delta t \nu_\chi\right) + \Delta t\left(\rho\sigma_\xi\sigma_\chi - \Delta_\chi \nu_\xi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}$$

$$p_{d|d} = \frac{\Delta_\xi\left(\Delta_\chi + \Delta t \nu_\chi\right) - \Delta t\left(\Delta_\chi \nu_\xi + \rho\sigma_\xi\sigma_\chi\right)}{2\Delta_\chi\left(\Delta_\xi + \Delta t \nu_\xi\right)}.$$

Each of these is censored in the approximation when the probabilities are invalid, that is,

$$p \equiv \begin{cases} p & if \;\; 0 \le p \le 1 \\ 0 & if \;\; p \le 0 \\ 1 & if \;\; 1 \le p \end{cases} \qquad \forall p \quad \left(p_{u|u}, p_{u|d}, p_{d|u}, p_{d|d}\right).$$

We see that $\lim\limits_{\Delta t \to 0} q_t = \frac{1}{2}$ for each conditional probability ($p_{u|u}$, $p_{d|u}$, $p_{u|d}$, $p_{d|d}$), which means that in the limit, we have convergence to the instantaneous drift, so

$$\lim_{\Delta t \to 0} \mu_t(\chi,t) = \kappa\chi - \rho\sigma_\chi = \mu(\chi,t).$$

The local variance is equal to the instantaneous variance, since there is no dependence on Δt, therefore we also satisfy the final condition

$$\lim_{\Delta t \to 0} \sigma_t^2(\chi,t) = \sigma^2(\chi,t).$$

3.3 TWO CORRELATED ONE-FACTOR MEAN-REVERTING MODELS

Another goal of this research was to develop a discrete binomial representation of two one-factor mean-reverting diffusions, again using a similar approach to the bivariate binomial approximations of Boyle (1988). The same basic principles used in developing the discrete two-factor commodity price model apply, but in this case both of the individual processes are mean-reverting and therefore approximated using the Nelson and Ramaswamy approach.

Both processes to be modeled are assumed to follow simple one-factor Ornstein-Uhlenbeck processes. The price process for commodity one is $dX_t = \kappa(\overline{X} - X_t)dt + \sigma_X dz_X$ and the price process for commodity two is $dY_t = \kappa(\overline{Y} - Y_t)dt + \sigma_Y dz_Y$. The relationship between the increments of the two processes is given by $dz_X dz_Y = \rho_{XY} dt$.

A two-dimensional binomial approximation can be developed for this process, which results in a four-branch chance node for each discrete period, as before. The probabilities for the joint process can be derived using the same steps as above, but with the drift of the processes as $\nu_X = \kappa(\overline{X} - X_t)$ and $\nu_Y = \kappa(\overline{Y} - Y_t)$, respectively. The equal up and down jump sizes for each process are $\Delta_X = \sigma_X \sqrt{\Delta t}$ and $\Delta_Y = \sigma_Y \sqrt{\Delta t}$, respectively. As with the short term-long term model, at some nodes it may be necessary to censor probabilities when the degree of mean reversion required from a particular state results in probabilities greater than one (upward force of reversion) or less than zero (downward force of reversion). Bayes' Rule can again be applied to split the joint distribution into a marginal distribution for one of the commodities and a conditional distribution for the other. The difference in this case is that both the marginal and joint probabilities could be censored. The marginal distribution will be of the same form as that for the one factor Ornstein-Uhlenbeck process shown earlier in this section.

As shown in the following schematic, the first node is a binomial node for the marginal price process for commodity X, followed by a binomial node for the conditional price process for commodity Y.

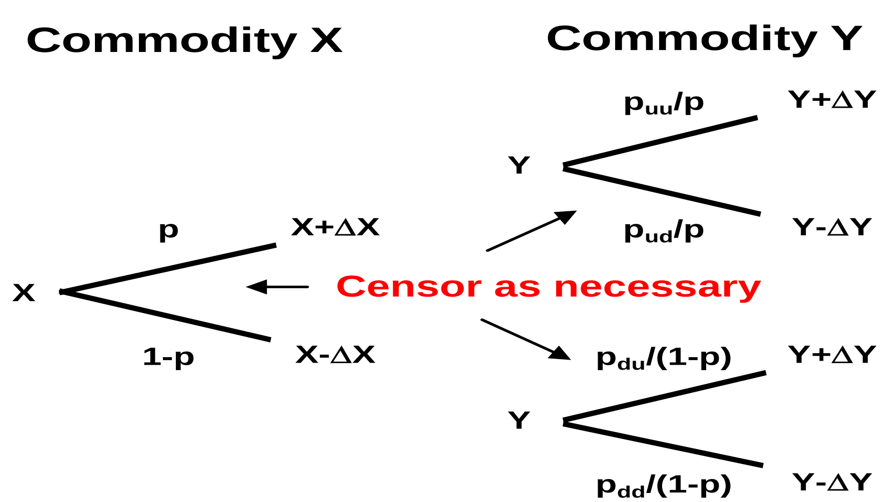

Figure 3.3 – Splitting the Four-branch node into Marginal and Conditional Steps

As with the two-factor model, convergence must be shown for both approximations. The proof of convergence for the binomial approximation of a one-factor Ornstein-Uhlenbeck process was given in Nelson and Ramaswamy (1990). For the conditional Ornstein-Uhlenbeck process the same steps that were used with the two-factor model earlier in this section are used, with the only difference being the drift term for commodity X , which changes from the GBM form, $\nu = \mu - \frac{\sigma^2}{2}$, to the form of $\nu = \kappa(\overline{X} - X_t)$. It is straightforward to see that this also meets the convergence criteria.

4. NUMERICAL RESULTS

The modeling approach based on Nelson and Ramaswamy's approximation can be applied to mean-reverting processes to facilitate the evaluation of real options on commodity price-contingent projects. As discussed in Section 2, this can either be done in binomial lattice or binomial tree format, and both will be demonstrated in this Section.

4.1 ONE-FACTOR MEAN-REVERTING MODELS

The first example is for a one-factor Ornstein-Uhlenbeck process. This is implemented in decision tree format with example parameters for a hypothetical process for oil price as follows: beginning price $Y_0 =$ ln($20), mean reversion coefficient $\kappa =$ 0.4, process volatility $\sigma =$ 0.2, and long-term mean price $\overline{Y} =$ ln($25). In practice, these parameters could be obtained from historical data. The objective here is to illustrate the approach by modeling prices over three years, beginning with a initial partition into three annual periods, so that $\Delta t = \frac{3}{3} = 1$.

A solved decision tree, which shows the endpoint values, probabilities, and expected value for price in the third period, is shown in Figure 4.1. It is evident upon inspection that values are recurring in the tree, as would be expected for a tree representation of a recombining binomial lattice. It is also evident from this figure that the probabilities are calculated at each node to reflect the mean reversion, and that there is one case where the probabilities are censored (after two down moves) due to the upward force of reversion.

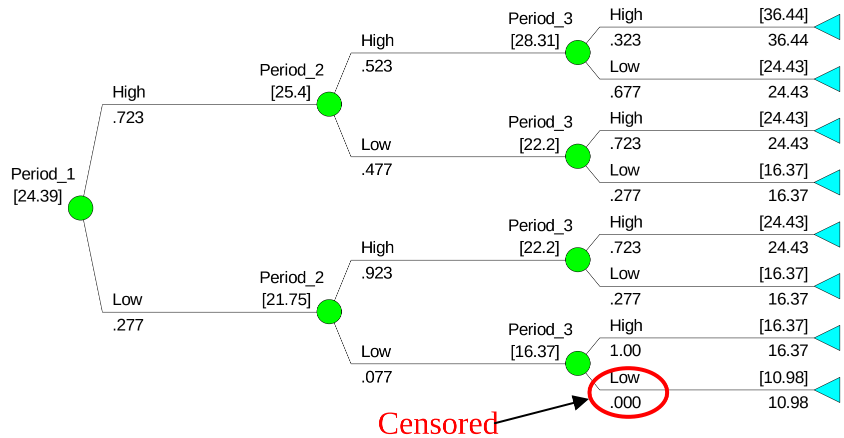

Figure 4.1 – Solved Three-Period Decision Tree for Third Period Price

To investigate convergence with this tree, the length of the time period can be reduced in several successive increments. In this case the three-year time horizon is divided up into an increasing number of steps, according to the following sequence: $\Delta t = \dfrac{3}{6} = 0.5$, $\Delta t = \dfrac{3}{12} = 0.25$, $\Delta t = \dfrac{3}{18} = 0.1667$, and $\Delta t = \dfrac{3}{24} = 0.125$. As an example, the tree for $\Delta t = 0.25$ is solved in Figure 4.2.

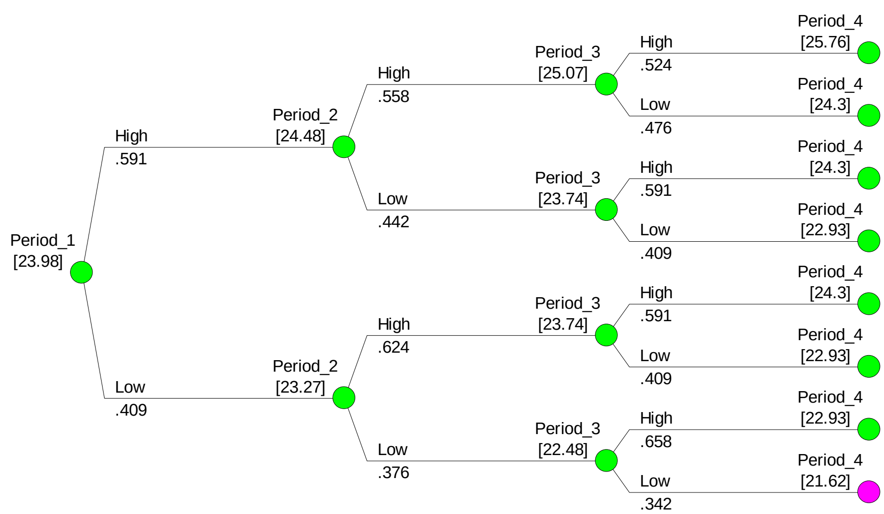

Figure 4.2 – Solved 12-Period Decision Tree for Twelfth Period Price (Partial View)

The price to which these models should converge can be calculated using the following expression for discretized values from the Ornstein-Uhlenbeck process:

$$Y_t = Y_{t-1} e^{-\kappa \Delta t} + \overline{Y}(1 - e^{-\kappa \Delta t})$$

Using the parameters for this example, the oil prices for year 1, year 2, and year 3 are \$21.83, \$23.07, and \$23.91, respectively. The convergence of the binomial model to these values is shown in Figure 4.3.

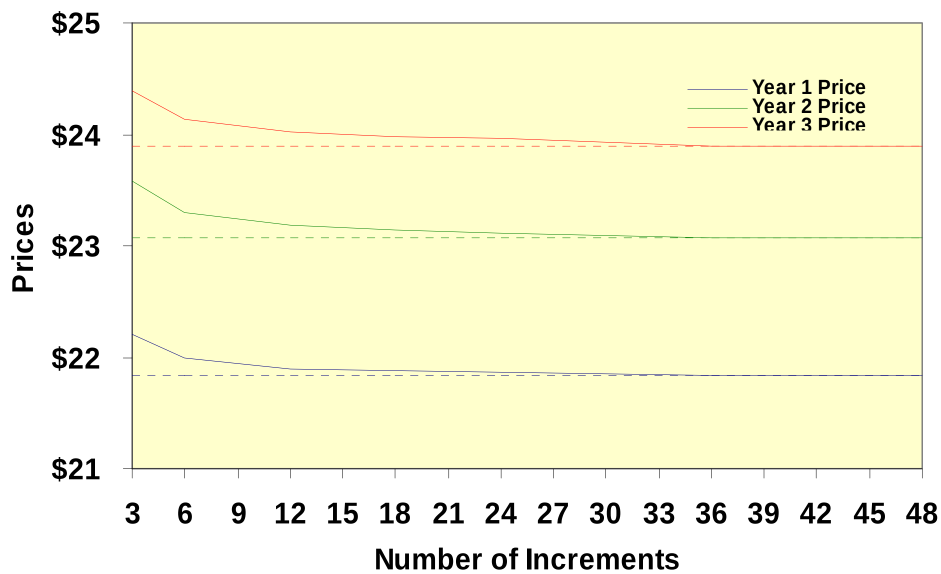

Figure 4.3 – Convergence of Prices for Years 1, 2, and 3

The initial values are biased upward by ~2% as a result of the censorship of low values at the very bottom of the diffusion, which were seen in Figure 4.1. However, the approximation converges rapidly, yielding values within 1% of true values if we use half-year periods ($\Delta t = \frac{3}{6}$) and values within 0.5% by using quarter-year periods ($\Delta t = \frac{3}{12}$).

While this implementation in decision tree format shows the model to be converging, for practical computational times it is limited to about 30 time steps, as the number of endpoints in the tree grows rapidly (2^{30}, or 1.07×10^{9} endpoints for 30 steps). In this case, a more efficient lattice-based algorithm in Visual Basic or other programming language can be used. Using this approach and further decreasing the period length to $\Delta t = \frac{3}{48}$ yields the convergence behavior shown in Figure 4.3.

4.2 TWO-FACTOR MEAN-REVERTING MODELS

The next objective is to implement the approach that was developed in the previous section for tree or lattice construction for a two-factor model and to test its convergence. To accomplish this, the Schwartz and Smith (2000) two-factor model is first implemented in decision tree format, using parameter data estimated by Schwartz and Smith from oil price data from 1/2/90 to 2/17/95. Based on this data, the current spot price of oil is \$19.61 and the parameters for the model are: $\kappa =$ 1.49, $\sigma_\chi =$ 28.6%, $\xi_0 =$ ln(\$17.41) = 2.857, $\mu_\xi =$ 1.6%, $\sigma_\xi =$ 14.5%, and $\rho =$ 0.3. The example of modeling prices over a three year period is used again, and the initial partition of time is into three annual periods ($\Delta t = \frac{3}{3} =$ 1). A decision tree model for the price of oil in the third period is shown in Figure 4.4. Nodes denoted with "Mean *i*" in the figure are the binomial nodes for the long-term mean ξ for period *i*, and nodes denoted with "X*i*" are the binomial nodes for the conditional process for the short-term deviation χ for period *i*.

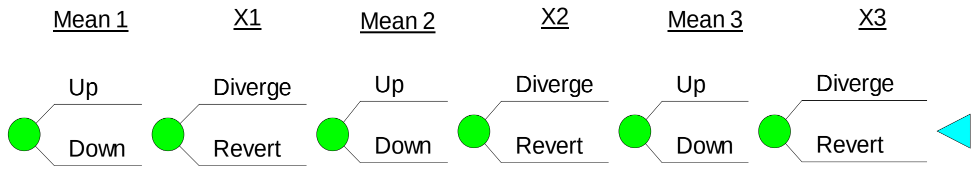

Figure 4.4 – Simple Three-Period Decision Tree – Two-Factor Model

The solved tree, with some nodes collapsed to save space, is shown Figure 4.5.

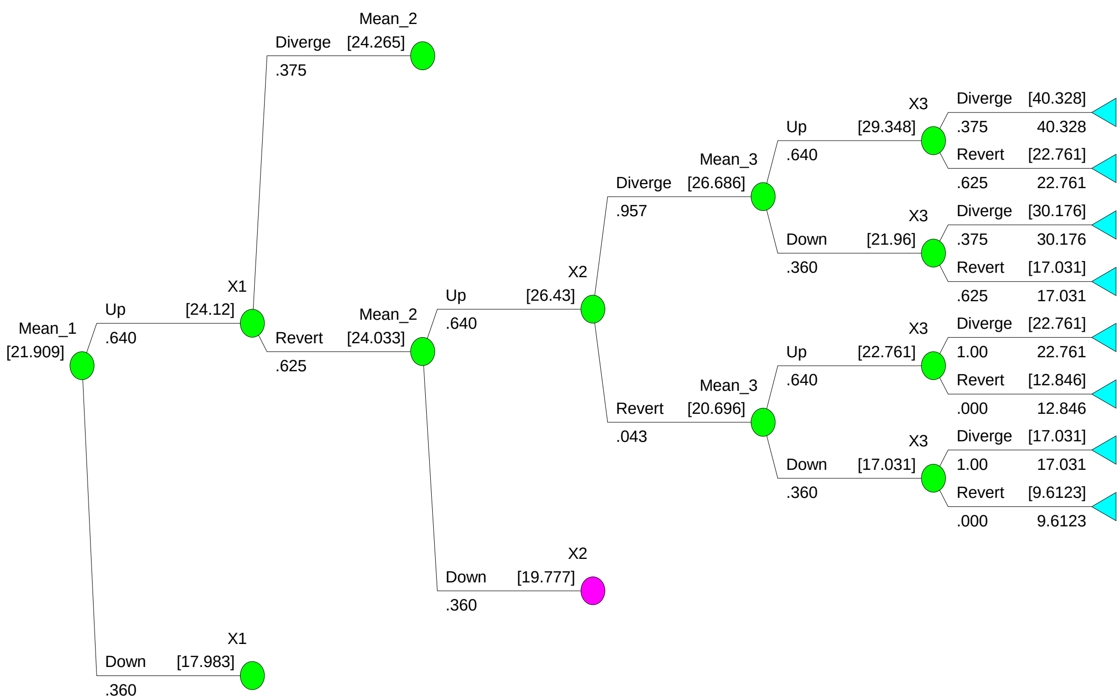

Figure 4.5 – Solved Decision Tree for Third Period Price (Partial View)

It is again evident from the recurring terminal values of the one expanded path that this process can be modeled as a binomial tree. Figure 4.5 also shows that the probabilities in the nodes for the long-term mean are constant, following GBM diffusion, while the probabilities in the short-term deviation nodes change to reflect the required degree of mean reversion.

To again investigate convergence of the approximation numerically, the length of time period is reduced in several increments. The values with large time increments again exhibit significant ($\approx \pm 5\%$.) error, but convergence to values within 1% are achieved rapidly (by $\Delta t = \frac{1}{6}$). As was the case with the one-factor model, we can continue to reduce the period length to more fully investigate convergence behavior by switching to a lattice-based algorithm, and we next discuss how to accomplish this.

4.2 LATTICE-BASED IMPLEMENTATION FOR TWO-FACTOR MODELS

Implementation of the approach developed in the Section 3 is fairly straightforward in decision tree format, but requires a few additional steps in lattice format. Although the endpoints of the binomial nodes are recombining, when there are two separate factors the procedure for capitalizing on the recurring values changes. To provide an example of where these values occur in an expanded tree, and how they can be arranged for a two dimensional lattice, the first couple of steps are shown explicitly in Figure 4.6. From this figure, in which the same parameter values as above are used, it is evident that while there are sixteen endpoints at the end of the second period, there are only nine unique values.

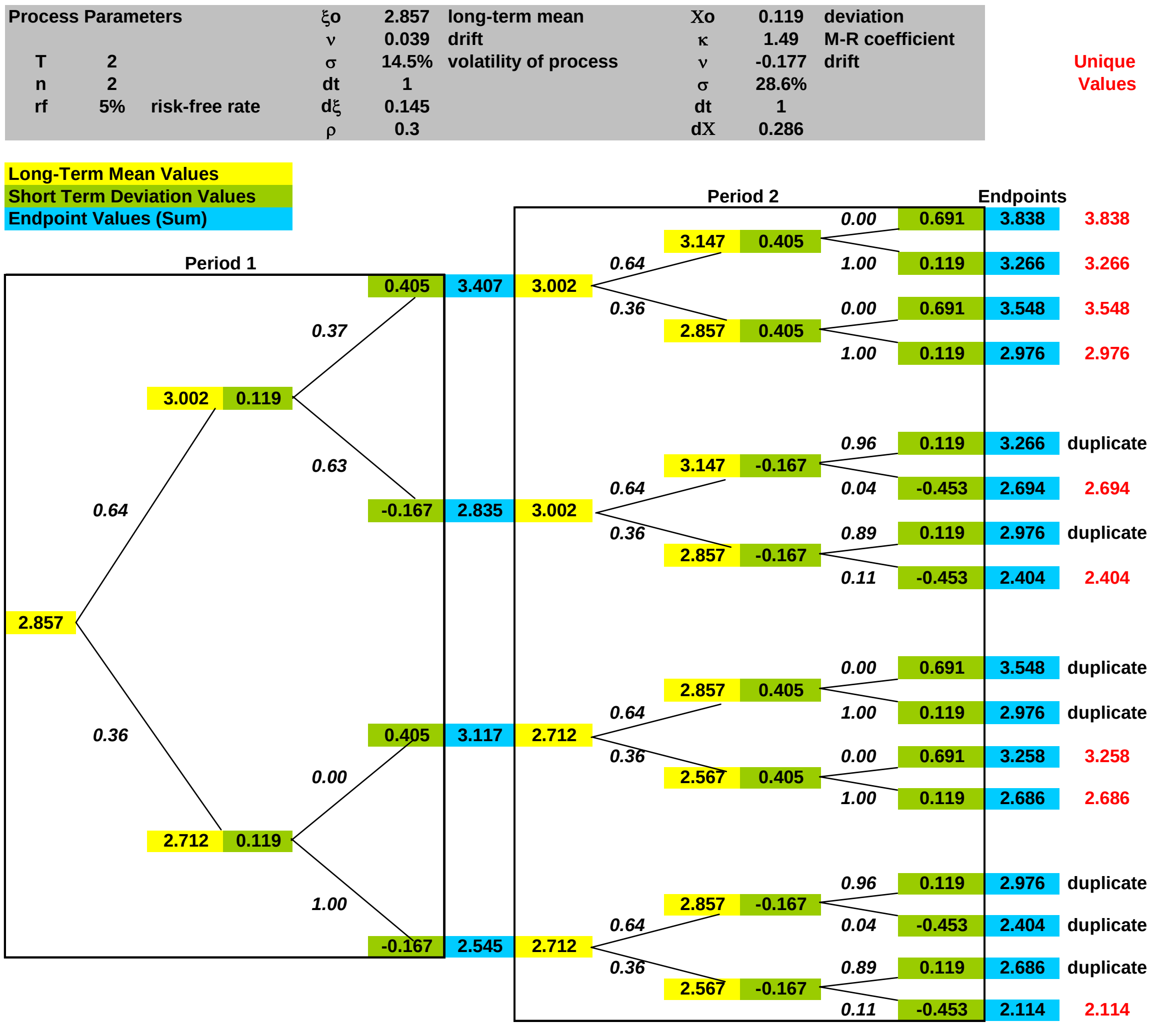

Figure 4.6 – Distribution of Unique Endpoints after Two Periods

The distribution of the endpoints appears to be arbitrary based on this one-dimensional view, however as discussed in Clewlow and Strickland (2000) this can be shown in two dimensions to lend more intuition toward the development of an algorithm to build the lattice. To show this, the same example shown above is represented in this alternative representation in Figure 4.7. In this view, shifts in the mean are indicated by the column position and shifts in the deviation from the mean are indicated by the row position.

Period 0

X					
0.691					
0.405					
0.119			**2.976**		
-0.167					
-0.453					
ξ	**2.567**	**2.712**	**2.857**	**3.002**	**3.147**

Period 1

X					
0.691					
0.405		**3.117**		**3.407**	
0.119					
-0.167		**2.545**		**2.835**	
-0.453					
ξ	**2.567**	**2.712**	**2.857**	**3.002**	**3.147**

Period 2

X					
0.691	**3.258**		**3.548**		**3.838**
0.405					
0.119	**2.686**		**2.976**		**3.266**
-0.167					
-0.453	**2.114**		**2.404**		**2.694**
ξ	**2.567**	**2.712**	**2.857**	**3.002**	**3.147**

Figure 4.7 – Distribution of Endpoints in Two Dimensions

The following figure shows the relationship between a standard branching representation and the two-dimensional view.

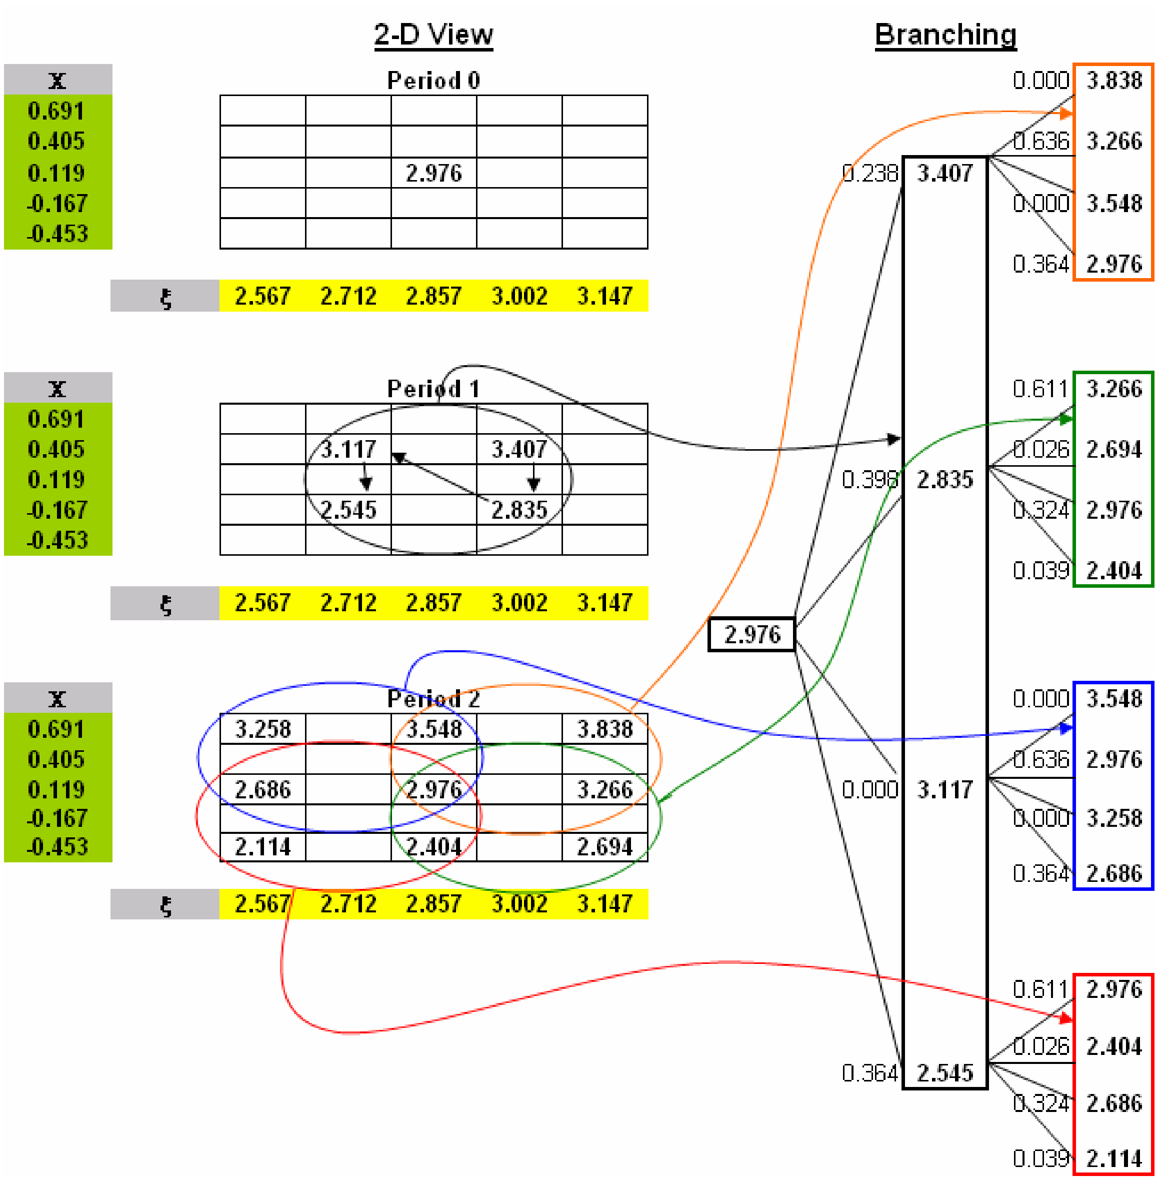

Figure 4.8 – Progression of Endpoints in Two Dimensions

With this approach, an algorithm can be developed to build a lattice in a two-dimensional array at each step. An example of such an algorithm is presented in Clewlow and Strickland (2000) for the case of two GBM's and based on four branch nodes at each step. In this case, one of the GBM's is replaced with a one-factor Ornstein-Uhlenbeck process and the four-branch node has also been replaced with a two node marginal-conditional sequence. However, it is relatively straightforward to convert the two node sequence back to a four branch node after testing to see whether censoring is required or not. Figure 4.9 shows and example of this process.

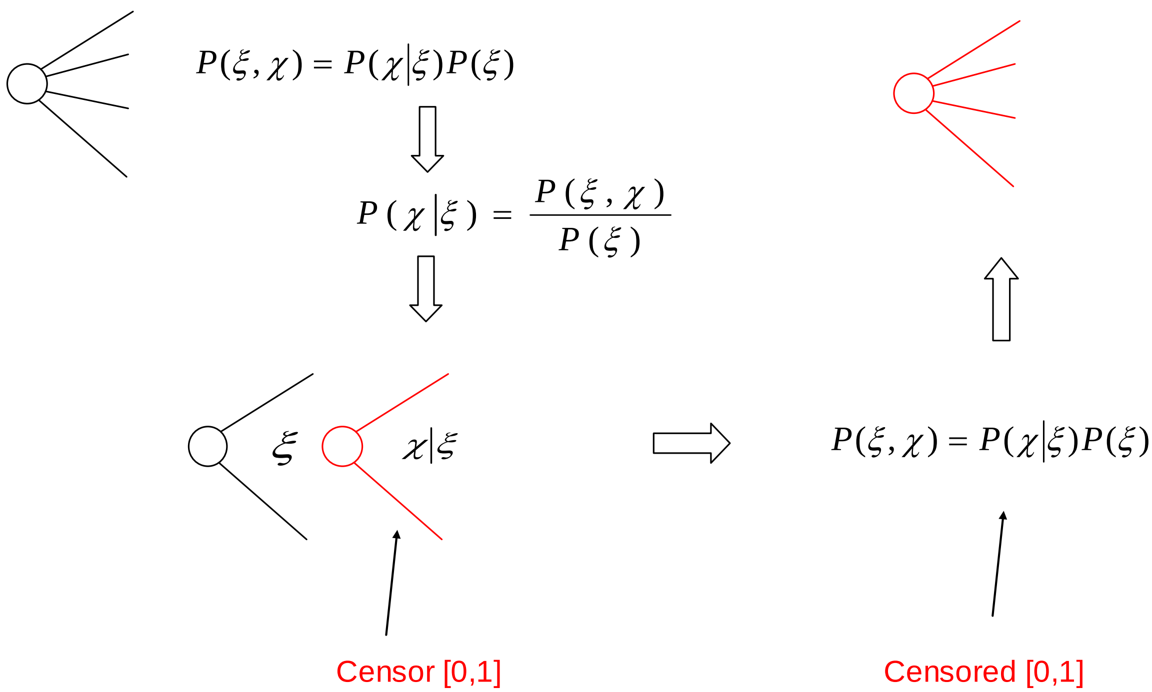

Figure 4.9 – Bayes Transformation and Inverse

By continuing past the number of feasible steps for a recursive decision tree algorithm, prices converge to the expected prices of $18.73, $19.22, and $20.06 for year 1, year 2, and year 3, respectively, as shown in Figure 4.10.

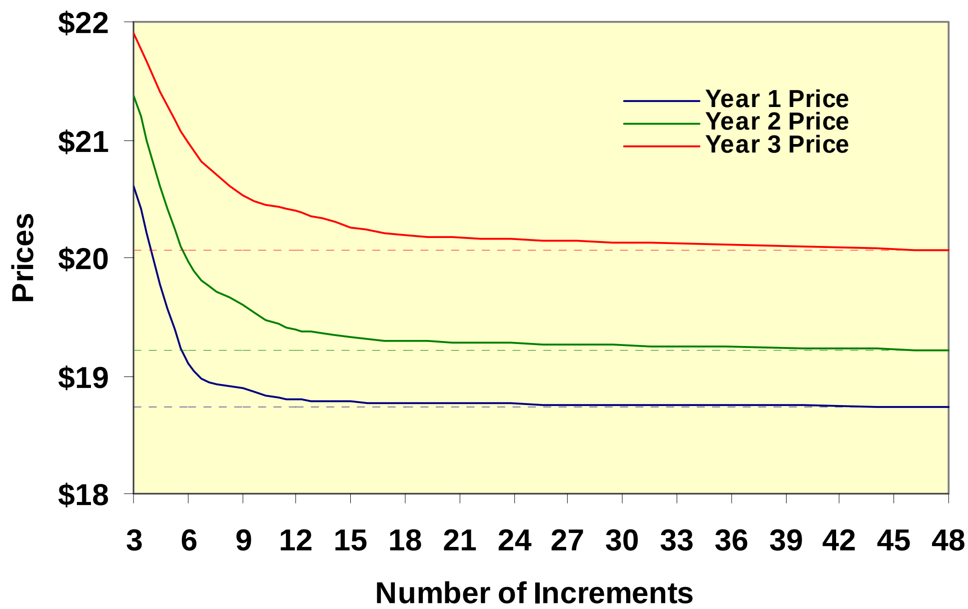

Figure 4.10 – Convergence of Two-Factor Prices

The figure also shows that convergence is slower for prices in periods farther out in time, which is to be expected as the errors in the model are compounded. Fortunately, the impact of such errors would be diminished by discounting in a valuation problem. In any case, for this example convergence to within what might be considered reasonable tolerance for a real option problem was achieved within the range of capabilities of recursive decision tree algorithms.

4.3 APPLICATION TO REAL OPTIONS

Thus far it has only been shown that this discrete approximation numerically converges to the expected prices. This is important, but does not ensure that option values calculated using the discrete approximation will converge to option prices that would be calculated from a continuous distribution. Convergence in distribution can be numerically tested by valuing a simple real option example and validating it against the results from existing approaches.

To illustrate such a test, the following example (from Hull, 1999) of an option associated with an oil project can be used. In this example, an oil producing firm is considering investing in a project that will deliver 2 million barrels of oil annually over a three-year period. The initial capital expenditure for the project is \$15MM, with annual fixed costs of \$6MM, and variable production costs of \$17/barrel. The risk-free discount rate is given to be 10%. Therefore, the expected net present value of the cash flows from the project can be calculated using the following formula:

$$NPV = -15 + ((P_1 - 17)*2 - 6)*e^{-0.1*1} + ((P_2 - 17)*2 - 6)*e^{-0.1*2} + ((P_3 - 17)*2 - 6)*e^{-0.1*3}$$

where P_1, P_2, and P_3 are future spot oil prices. Using the expected futures prices given in the example ($P_1 = \$22.00$, \$ $P_2 = \$23.00$, and $P_3 = \$24.00$) yields a net present value of -\$0.54MM.

If the model is calibrated to actual futures prices, we have a risk-neutral forecast of future oil prices. Risk-neutrality implies that the owner of this project could arrange for a hedge against the production using financial instruments in the commodities markets, and thereby guarantee the net present value that is derived when the futures prices are assumed as above. This of course assumes that there is no uncertainty in the amount of oil to be produced, or that the firm has a sufficient number of project of this type so that if this project under-delivers, some other project in the firm's portfolio will compensate, and vice-versa. The key advantage of having a risk-neutral forecast is that future cash flows resulting from the project's options can be discounted at the risk-free rate to obtain a valid project value. Values calculated in this manner reflect a consensus view of what the project would be worth, without having to consider the risk level of the project. Without a risk-neutral forecast, we would have no view on the relative risks associated with the project's payoffs, and therefore could only obtain a valuation based on an arbitrary discount rate.

Like many real projects, this example could have embedded options due to project managerial flexibility. For example, there may be an option to abandon the project with zero salvage value at different points during the project, with the following payoffs:

$$NPV = -15 \qquad \text{(during first year)}$$

$$NPV = -15 + ((P_1 - 17) * 2 - 6) * e^{-0.1*1} \qquad \text{(during second year)}$$

$$NPV = -15 + ((P_1 - 17) * 2 - 6) * e^{-0.1*1} + ((P_2 - 17) * 2 - 6) * e^{-0.1*2} \qquad \text{(third year)}$$

The ability to exercise this option to avoid bad outcomes (negative cash flows) if prices fall below a certain level changes the riskiness of this project. It also obviously changes the value relative to the deterministic expected value case.

Hull (1999) uses a discrete trinomial (three-branch) tree to solve this example. Compared to a binomial tree, the extra branch gives an added degree of freedom to accommodate mean reversion. In Hull's approach, a trinomial tree is constructed to model a simple mean-reverting process and then calibrated so that the expected values match the given futures price in each period. Then decision nodes are added to the tree to reflect the abandonment decisions. The solved tree for this example is shown in Figure 4.11.

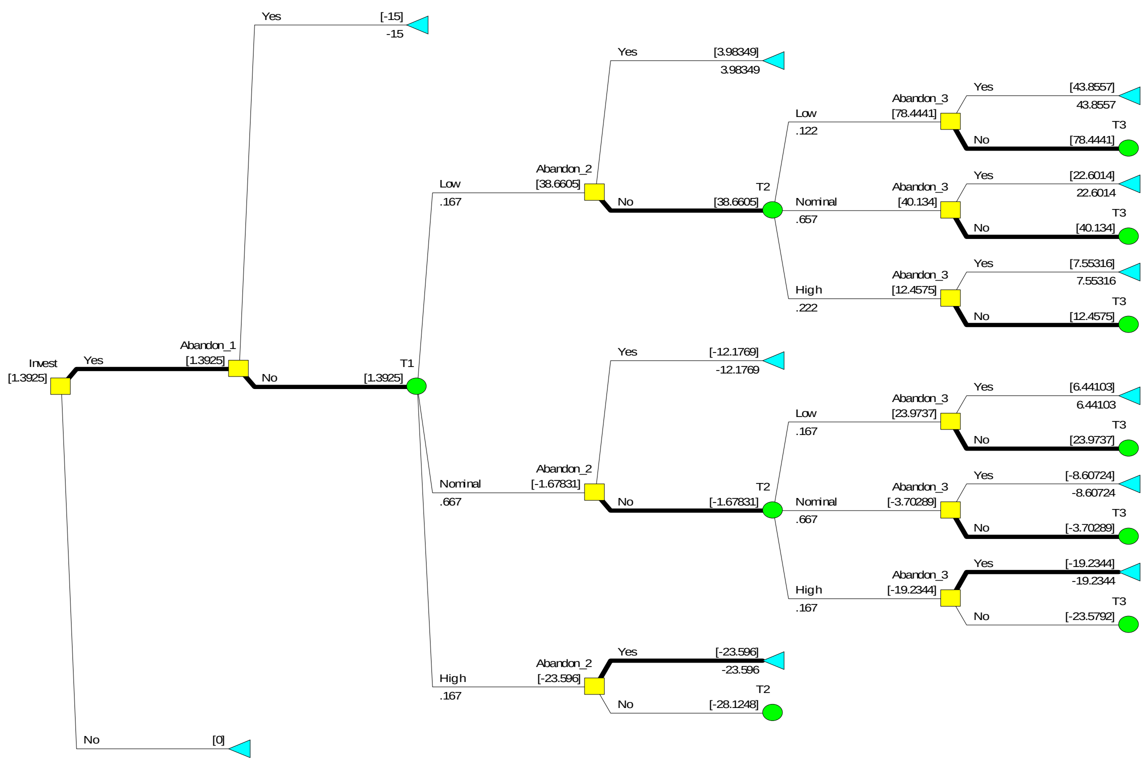

Figure 4.11 – Trinomial Tree for Valuing Abandonment Option

As the figure shows, we obtain a value of approximately \$1.40MM for the project with the option to abandon, so the incremental value of the option itself is \$1.94MM. If future oil prices were accurately modeled as a GBM, valuation of this option, which is analogous to an American Put, could be accomplished through straightforward application of a binomial model. However, a GBM model does not capture the mean reversion evident in the term structure of futures prices. As a result, if we fit a GBM to the futures data and solve with a binomial model, we obtain an option value of $3.60MM.

To verify Hull's results, we can also work the example using a simulation-based approach. Simulation does not require discretization of outcomes at each increment, and can easily model virtually any form of stochastic process, however some method of evaluating the decisions at each instance must be derived. Using the Longstaff and Schwartz (2001) approach discussed earlier, and implementing with the simple one-factor mean-reverting process used by Hull, we can find a value in agreement with the result from the trinomial tree method.

We next switch to the Schwartz and Smith two-factor model and fit it to the futures prices from the example by finding the parameters that minimize the squared deviations between predicted and actual futures prices. Under this approach, the parameters are: $\xi_0 =$ 3.374, $\chi_0 =$ -0.378, $\kappa =$ 0.3, $\sigma_\chi =$ 15%, $\mu_\xi =$ 2%, $\sigma_\xi =$ 12%, and $\rho =$ 0.3. Using this process in the Longstaff and Schwartz simulation approach, we again arrive at the same result of $1.94 for the value of the abandon option.

Finally, we use our binomial approximation of the two-factor model to work the example and validate against the above results. We use the decision tree implementation, as previously shown in Figure 8, and add decision nodes to reflect the option value and to convert the terminal payoffs from prices to project cash flows. Carrying out these steps results in the tree shown in Figure 4.12, and the solution is shown in Figure 4.13.

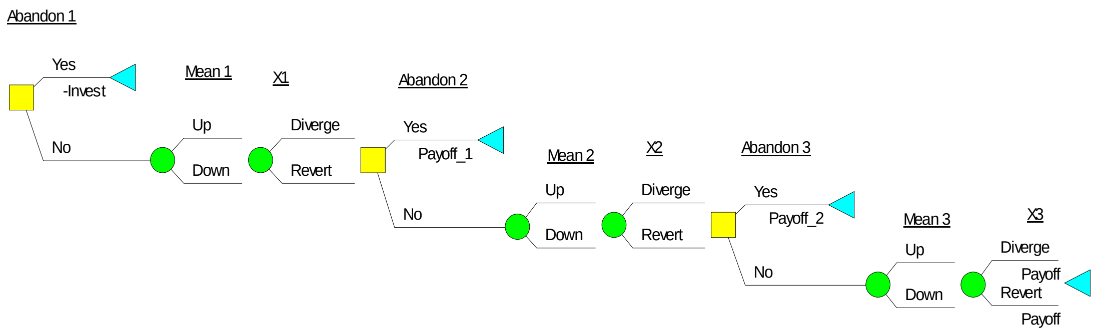

Figure 4.12 – Decision Tree for Valuing Abandonment Option

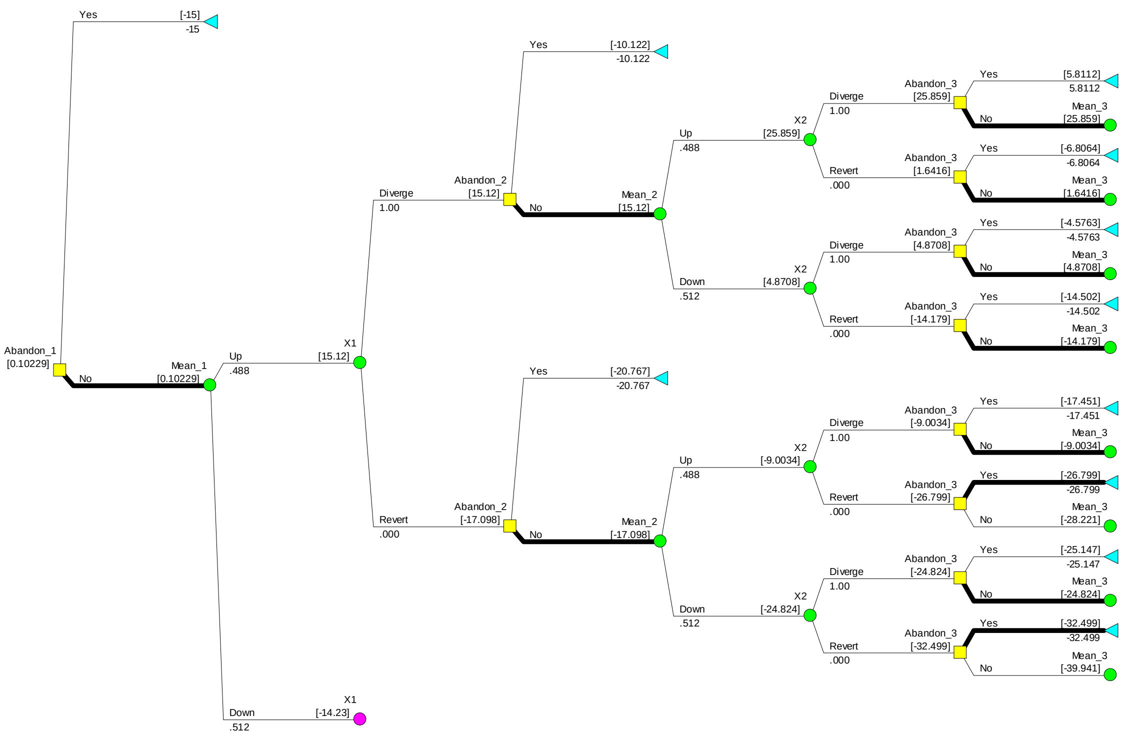

Figure 4.13 – Solved Decision Tree for Valuing Abandonment Option (Partial View)

The solution for project value with the abandonment option is $0.10MM, which yields an incremental option value of $0.64MM. As shown by Amin and Khanna (1994) for the case of an early exercise option such as the abandon option in this example, if we have convergence in distribution, then we should have convergence in option value as well. We show that this is indeed the case by again reducing the time increments, this time between calculations of project payoffs. The results are shown in Figure 4.14.

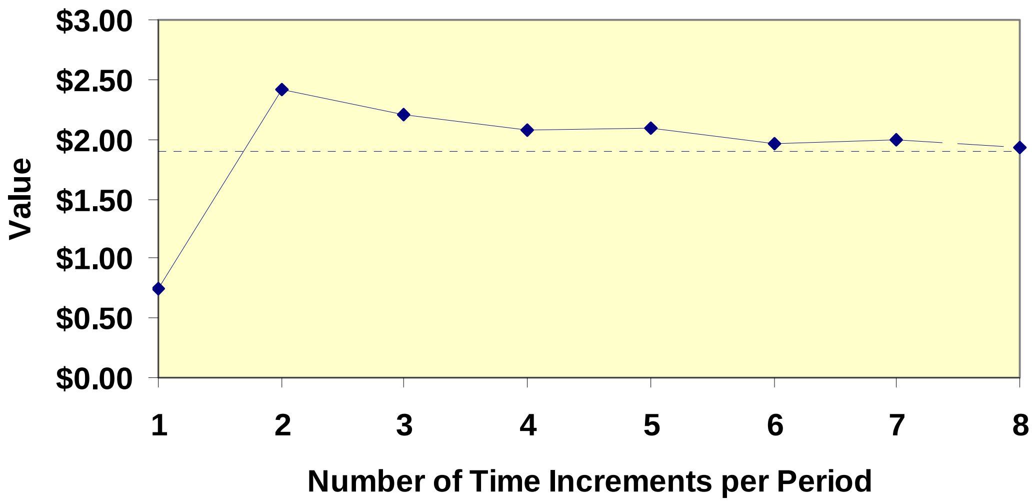

Figure 4.14 – Convergence in Option Value

The fact that this approach converges to the same solutions for a simple option as the ones provided by a Monte Carlo simulation approach and a trinomial tree is important. However, this approach is much more flexible than the simulation-based approach, and provides a simple one-step method for valuing multiple concurrent options with complex payoff characteristics. Once the tree for underlying asset value is created, this approach can accommodate most any combination of options and payoff characteristics simply by adding decision nodes and terminal payoff statements at the appropriate locations in the tree. It can also easily accommodate two-factor processes with superior out-of-sample performance relative to the Hull one-factor process. This is important if, for example, the duration of the project extends past the available future data, as is often the case in real option valuation

5. PARAMETER ESTIMATION FOR A TWO-FACTOR MODEL

For valuation problems that have underlying stochastic processes that are mean-reverting, such as projects with cash flows driven by mean-reverting commodity prices, a necessary first step is constructing a model of the problem is to select an appropriate diffusion model and then determine its parameters empirically from data. Although futures markets are primarily in existence to provide an inexpensive way to transfer risks, a side benefit of these markets is that they also impound information about commodity prices that can be used to specify a diffusion model.

5.1 DIFFUSION MODELS

In some cases, a simple one-factor model may be appropriate for modeling commodity price evolution; in others it may be necessary to utilize a diffusion that incorporates more than one factor to model more complex interactions, such as economic supply and demand effects. Examples of this type of model include Gibson and Schwartz (1990), Schwartz (1997), Schwartz and Smith (2000), and others, as was discussed briefly in Sections 2 and 3. Schwartz (1997) compared a one-factor model, a two-factor model that incorporates mean-reverting convenience yield as the second factor, and a three-factor model that adds stochastic interest rates to the two-factor model in terms of each model's ability to fit commodity futures prices. That study found that two- and three-factor models outperformed one-factor models, and the three-factor model produced only marginally better results relative to the two-factor model. Hilliard and Reis (1998) investigated the differences in results for similar one-, two-, and three-factor commodity price models for valuing both financial and real assets. They found significant differences between one- and two-factor models when very high or very low convenience yields occur during the term, and that the difference between the two- and three-factor models depends on the interrelationship between interest rate volatility, correlation between the spot price and interest rate, and correlation between the convenience yield and interest rate.

In this section two-factor models are discussed in more detail and the various approaches to estimating model parameters are tested on empirical data and contrasted. Although we make some comparisons based on statistics, we note that this is not an exhaustive, scientifically valid comparison of these methods. There are two basic types of two-factor models. The first approach, used by Gibson and Schwartz (1990), Schwartz (1997), and Ribeiro and Hodges (2004) is to model price as a GBM as the first factor, and nest within the drift function of the price process a mean-reverting process for convenience yield. Hull and White (1994b) also use a variation of this approach in their two-factor model; however they use their fitted mean-reverting formulation as the process for the first factor instead of a GBM. The second approach is to decompose price into factors for the long-term mean, which is specified with a GBM process, and the short-term deviation from the long-term mean, which is modeled as a one-factor mean-reverting process. Schwartz and Smith's (2000) short-term/long-term model is the primary example of this approach. As mentioned in Section 4, this approach is more computationally convenient from the perspective of the discrete modeler, because the two factors are connected only by the correlation of their increments.

The basic rationale for the Schwartz and Smith (2000) model is to draw on the valid arguments of both primary single-factor models. Prices at a basic level should be expected to grow at a constant rate over time with variance increasing in proportion to time, which is behavior that can be modeled with a GBM. In the short term, however prices will also be affected by supply and demand conditions. Since these effects are short-lived, they would be expected to go away over time, which can be modeled with a process reverting to a mean of zero. The Schwartz and Smith model accommodates both types of behavior by introducing a bifurcation of the time horizon.

Following the nomenclature of Schwartz and Smith (2000) the long-term equilibrium price and deviation from the equilibrium price at any point are denoted as ξ_t and χ_t, respectively. As discussed in Section 3, the price is the sum of the two factors:

$$Y_t = e^{\chi_t + \xi_t},$$

where the two processes are:

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \qquad \text{(GBM for long-term mean price)}$$

$$d\chi_t = \kappa(0 - \chi_t)dt + \sigma_\chi dz_\chi \qquad \text{(Mean-reverting process for the deviation)}$$

and the increments of the two processes are correlated:

$$dz_\xi dz_\chi = \rho_{\xi\chi} dt.$$

This formulation indicates that there are five parameters required to specify this model: $\mu_\xi, \sigma_\xi, \kappa, \sigma_\chi,$ and $\rho_{\chi\xi}$.

Since the state variables ξ_t and χ_t are unobservable, we need some way to link them to observable information to in order to determine the five parameters above that define their stochastic process. We can use futures price data for that purpose.

Under risk-neutral valuation, the Schwartz and Smith process should be transformed with the addition of two parameters, λ_χ and λ_ξ, to adjust the drift of each process to produce a risk-neutral price model. The function of these two parameters, called the short-term deviation risk premium and equilibrium risk premium respectively, is to transform the two processes so that cash flows generated from the model can be discounted at the risk-free rate. The resultant formulation for the two-factor process then becomes:

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dz_\xi^*$$

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\xi^*$$

Denoting the risk-neutral drift as $\mu_\xi^* = \mu_\xi - \lambda_\xi$, the first equation can be written as:

$$d\xi_t = \mu_\xi^* dt + \sigma_\xi dz_\xi^*$$

The result is that the log of future spot price is normally distributed with the following revised mean and variance:

$$E\left[\ln(Y_t)\right] = e^{-\kappa t}\chi_0 + \xi_0 - (1 - e^{-\kappa t})\lambda_\chi / \kappa + \mu_\xi^* t$$

$$Var[\ln(Y_t)] = (1 - e^{-2\kappa t})\sigma_\chi^2 / 2\kappa + \sigma_\xi^2 T + 2(1 - \exp^{-\kappa t})\rho_{\chi\xi}\sigma_\chi\sigma_\xi / \kappa$$

A complete derivation of these formulas can be found in Schwartz and Smith (2000). Under risk-neutral valuation, the futures prices will equal the expected spot prices (Black, 1976). Therefore the expectation and variance can be used to derive the following expression for futures prices:

$$\ln(F_{T,0}) = e^{-\kappa T}\chi_0 + \xi_0 + A(T)$$

where,

$$A(T) = \mu_\xi^* T - (1-e^{-\kappa T})\frac{\lambda_\chi}{\kappa} + \frac{1}{2}\left((1-e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2(1-e^{-\kappa T})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right)$$

There are now seven parameters required to specify this model: $\kappa, \sigma_\chi, \sigma_\xi, \mu_\xi, \mu_\xi^*, \rho_{\chi\xi}$ and λ_χ, however we now have a method to link these parameters to observable data.

There are three primary methods for estimating these parameters, two of which use historical data and one of which implies estimates from forward-looking data. The first approach using historical data is Kalman filtering with maximum likelihood estimation of the parameters.

5.2 ESTIMATION USING THE KALMAN FILTER

The Kalman filter is a recursive procedure for estimating unobserved state variables based on observations that depend on these state variables (Kalman, 1960). In this case, the Kalman filter can be applied to estimate the unobservable state variables χ_t and ξ_t in the Schwartz and Smith model using the futures pricing equation shown in the previous section. It is then possible to calculate the likelihood of a set of observations given a particular set of parameters. By varying the parameters and re-running the Kalman filter, the parameters that maximize the likelihood function can be identified. A detailed description of this technique can be found in Harvey (1989).

For the Kalman filter, Schwartz and Smith (2000) specify the *transition equation* as:

$$x_t = c + Gx_{t-1} + \omega, \qquad t = 1,\ldots,n_T$$

where,

$x_t = [\chi_t, \xi_t]$ is a 2×1 vector of state variables

$c = [0, \mu_\xi \Delta_t]$ is a 2×1 vector

$G = \begin{bmatrix} e^{-\kappa\Delta t} & 0 \\ 0 & 1 \end{bmatrix}$ is a 2×2 vector of state variables

ω is a 2×1 vector of serially uncorrelated normally-distributed disturbances with:

$E[\omega_t] = 0$ and

$$Var[\omega_t] = W = Cov[\chi_{\Delta t}, \xi_{\Delta t}] = \begin{bmatrix} (1-e^{-2\kappa t})\dfrac{\sigma_\chi^2}{2\kappa} & (1-e^{-2\kappa t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1-e^{-2\kappa t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2 t \end{bmatrix}$$

Δt is the length of time steps

n_T is the number of time periods

The corresponding *measurement equation* is:

$$y_t = d_t + F_t^{'} x_t + v_t, \qquad t = 1,\ldots,n_T$$

where,

$y_t = [\ln(F_{T1}),...,\ln(F_{Tn})]$ is a $n \times 1$ vector of observed (log) futures prices with maturities $T_1, T_2,...,T_n$

$d_t = \lfloor A(T_1),..., A(T_n)\rfloor$ is a $n \times 1$ vector

$F_t = \lfloor e^{-\kappa T_1} 1,..., e^{-\kappa T_n} 1\rfloor$ is a $n \times 2$ matrix

v_t is a $n \times 1$ vector of serially uncorrelated normally-distributed disturbances (measurement errors) with $E[v_t] = 0$ and $Cov[v_t] = V$.

With these two equations and a set of observed futures prices for different maturities, the Kalman filter can be run recursively beginning with a prior distribution of the initial values of the state variables (χ_0, ξ_0). A multivariate normal with mean vector m_0 and covariance matrix C_0 is assumed.

In each subsequent period, the next observation y_t and the previous period's mean vector and covariance matrix are used to calculate the posterior mean vector and covariance matrix. The mean and covariance of the state variables are given by:

$$E[\chi_t, \xi_t] = m_t = a_t + A_t (y_t - f_t)$$

$$Var[\chi_t, \xi_t] = C_t = R_t - A_t Q_t A_t^{'}$$

where,

$a_t = c + Gm_{t-1}$ (mean of (χ_t, ξ_t) based on what is known at $t-1$)

$R_t = G_t C_{t-1} G_t^{'} + W$ (covariance of (χ_t, ξ_t) based on what is known at $t-1$)

$f_t = d_t + F_t^{'} a_t$ (mean of period *t* futures price based on what is known at $t-1$)

$Q_t = F_t^{'} R_t F_t + V$ (covariance of period *t* futures price based on what is known at $t-1$)

$A_t = R_t F_t Q_t^{-1}$ (correction to predicted state variables a_t based on the difference between the (log) observed prices observed at time *t*, y_t, and the predicted price vector at time *t*, f_t)

As described in Harvey (1989), Chapter 3.4, Kalman filtering facilitates calculation of the likelihood of a set of observations given a particular set of parameters. In this case there are seven model parameters to estimate $\left(\kappa, \sigma_\chi, \sigma_\xi, \mu_\xi, \mu_\xi^*, \rho_{\chi\xi}, \lambda_\chi\right)$, along with the terms in the covariance matrix for the measurement errors (V). This can be simplified with the common assumption that the errors are not correlated with each other, so that V is diagonal with elements $\left(s_1^2, ..., s_n^2\right)$, as in Schwartz (1997) and Schwartz and Smith (2000). The general form for the log-likelihood function for a joint normal distribution is:

$$\ln(L) = -\frac{1}{2}\sum_{t=1}^{T} \ln\left|F_t\right| - \frac{1}{2}\sum_{t=1}^{T}\left(y_t - \hat{y}_{t|t-1}\right)^{'} F_t^{-1}\left(y_t - \hat{y}_{t|t-1}\right) + const.$$

where,

$$F_t = E\left[\left(y_t - \hat{y}_{t|t-1}\right)\left(y_t - \hat{y}_{t|t-1}\right)^{'}\right]$$

Schwartz and Smith (2000) use the maxlik routine in Gauss to numerically determine the estimates of the above parameters for their two-factor process based on two different crude oil futures price data sets. The first set covers the period from 1990 - 1995 and has prices for contracts with 1-, 5-, 9-, 13-, and 17-month maturities. The second set covers a different period, 1993 -1996, and includes more longer-term contracts, with a spread of 2-, 5-, 8-, 12-, 18-, 24-, 36-, 60-, 84-, and 108-month maturities.

In this section we use a three-part MATLAB routine on a data set covering futures contract maturities of 1-, 3-, 5-, 9-, 13-, and 17 months and compare both in-period parameter estimates to those mentioned above, as well as estimates from a more current and extensive data set, which covers the period from 1990-2005 as shown in Figure 5.1

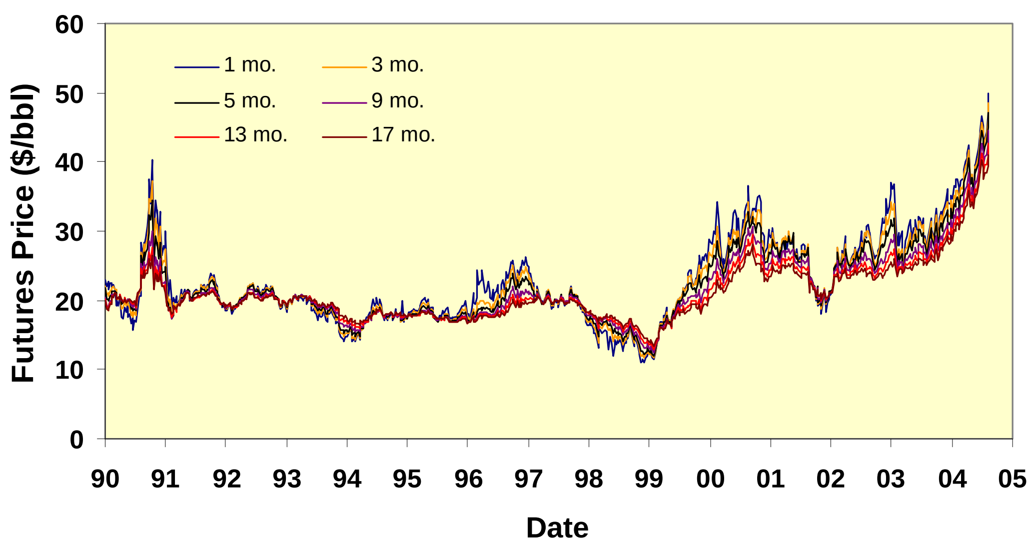

Figure 5.1 – Crude Oil Futures Data Set

The MATLAB routine, which is based on code developed by Jim Smith at Duke University, includes modules to read in and manipulate data sets, to return the likelihood function based on the Kalman filter, and to maximize the likelihood function. The first set of parameter estimates, covering the 1990 to 1995 period is shown in Table 5.1. Although the algorithm did not automatically find the parameters that globally maximize the likelihood function, by trying a few different starting parameter estimates, a solution was eventually found. Once this solution was determined, the starting parameter estimates could be changed within a limited range and the algorithm would still converge to the same solution. Thus, the algorithm did not always find a solution to the global maximization, but when it did, it was always the same solution.

		'90-'95		
Parameter		**S&S**	**Std. Err.**	**This Study**
Equilibribum drift rate	μ_ξ	-0.0125	0.0728	0.0116
Short-term mean-reversion rate	κ	1.4900	0.0300	1.5002
Short-term risk premium	λ_χ	0.1570	0.1440	0.2740
Short-term volatility	σ_χ	0.2860	0.0100	0.3411
Equlibrium volatility	σ_ξ	0.1450	0.0050	0.1623
Correlation in increments	$\rho_{\xi\chi}$	0.3000	0.0440	0.3519
Equilibrium risk-neutral drift rate	μ_ξ^*	0.0115	0.0013	0.0100
Standard deviation of error for Measurement Eq.	s_1	0.0420	0.0020	0.0408
	s_2	0.0060	0.0010	0.0028
	s_3	0.0030	0.0000	0.0042
	s_4	0.0000	0.0000	0.0019
	s_5	0.0040	0.0000	0.0051
	s_6	n/a		n/a

Table 5.1 – Comparison with Schwartz and Smith's '90 – '95 Results

The results shown above indicate good agreement between the algorithm used by Schwartz and Smith (2000) and the one used in this study. If we consider the confidence intervals around parameter estimates from Schwartz and Smith, only the estimates for short-term volatility and equilibrium volatility are outside the 95% intervals, however these two parameters also have very small standard errors. Table 5.2 shows the same information for the period from 1993 to 1996. Here we also see general agreement between the two different parameter estimates. Short-term mean-reversion rate, short-term volatility, and equilibrium risk-neutral drift rate fall outside the 95% confidence interval, owing again to the very low standard errors for those parameter estimates in the Schwartz and Smith study.

Parameter		'93-'96 S&S	Std. Err.	This Study
Equilibribum drift rate	μ_ξ	-0.0386	0.0728	-0.0554
Short-term mean-reversion rate	κ	1.1900	0.0300	1.5624
Short-term risk premium	λ_χ	0.0140	0.0820	0.0366
Short-term volatility	σ_χ	0.1580	0.0090	0.1912
Equlibrium volatility	σ_ξ	0.1150	0.0060	0.1026
Correlation in increments	$\rho_{\xi\chi}$	0.1890	0.0960	0.1721
Equilibrium risk-neutral drift rate	μ_ξ^*	0.0161	0.0012	0.0236
Standard deviation of error for Measurement Eq.	s_1	0.0270	0.0010	0.0422
	s_2	0.0060	0.0010	0.0086
	s_3	0.0000	0.0000	0.0000
	s_4	0.0020	0.0000	0.0027
	s_5	0.0000	0.0000	0.0000
	s_6	0.0050	0.0000	0.0041

Table 5.2 – Comparison with Schwartz and Smith's '93 – '96 Results

Figures 5.2 and 5.3 compare the forecasts that would be obtained in each case graphically. While our results duplicated most of those from Schwartz and Smith's, we might expect some differences due to slight differences in data. We did not have access to their sources, which were Knight-Ridder financial services (1990-1995) and Enron (1993-1996), and instead used data from Bloomberg. Given that global optimization of a function of seven variables is a challenging computational problem, slight differences between the estimates could also result from the use of different optimization routines. A Gauss optimization routine was used in the Schwartz and Smith study, whereas we used a MATLAB routine in this study.

By experimenting with changes to the different parameters, we found the forecasts were most sensitive to the equilibrium drift rate, equilibrium risk-neutral drift rate, and the short-term risk premium. The differences in results in Figures 5.2 and 5.3 are largely due to differences in the estimates of these parameters. In the case shown in Figure 5.2, the difference is almost entirely due to an estimate of the short-term risk premium that is 75% higher in our case than in that of Schwartz and Smith (2000). This may be due to slight differences in data sets, since even small differences in the period around the Gulf War from late 1990 through the first half of 1991 would have a significant impact on the estimated short-term risk premium. As shown in Figure 5.1, this was the period when the differences in prices for the different maturities were the most pronounced. We also note that the short-term risk premium was the most difficult to estimate, as it had the highest standard error of all the estimates.

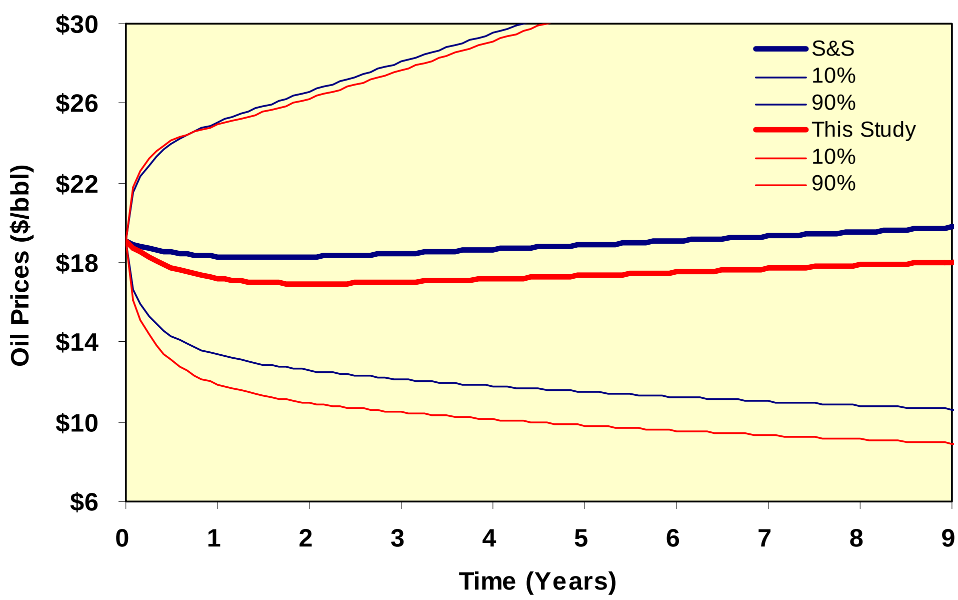

Figure 5.2 – Comparison of Forecasts Based on ’90 – ’95 Data

The short-term risk premium was also higher is the case shown in Figure 5.3, although not by as much as the previous case because the Gulf War period is not included. This difference was offset by a higher estimated equilibrium drift rate. In this case, the variation is most likely due to the slightly different maturities used in the data for this study and in the Enron data.

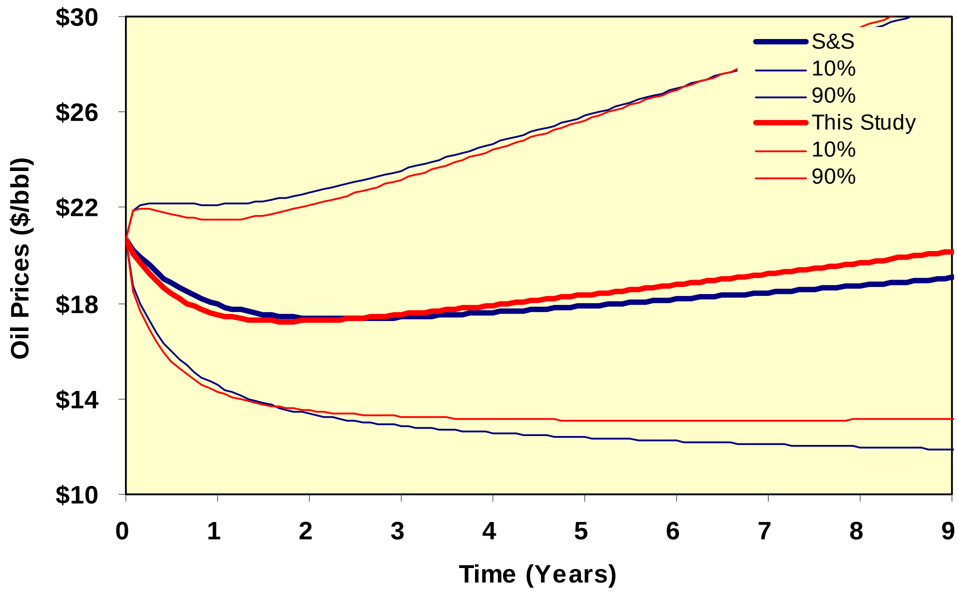

Figure 5.3 – Comparison of Forecasts Based on '93 – '96 Data

In general, however, the forecasts are similar in structure and standard errors for the parameter estimates were also similar to those obtained by Schwartz and Smith. These ranged from just over 0.0668 for the equilibrium drift rate estimate of 0.0116 for the 1990 to 1995 period, to 0.0056 for the equilibrium volatility of 0.1026 for the 1993 to 1996 period. Consequently, from the perspective of our study, we find that the parameter estimates from the Schwartz and Smith study that fall outside the 95% confidence intervals are nearly the same as those noted above from the reverse perspective. The only change is the addition of the short-term risk premium for the 1990 to 1995 period and the equilibrium volatility for the 1993 to 1996 periods as estimates falling outside the confidence intervals, due to smaller errors in our study.

We next fit the data over the expanded time horizon, from January 1990 through September 2004, with the results compared to the two prior fits shown in Table 5.3. The fit to the expanded data set picks up the events of 1990 – 1991, and the effect on the short-term risk premium is again evident. We also note that the equilibrium drift rate and risk-neutral equilibrium drift rate are the highest and lowest, respectively, of those estimated in any of the three cases, indicating that the long-term risk premium has increased. This is likely due to the run-up in prices since 2000, and the uncertainty about the long-term equilibrium level.

Parameter		'90 - '95	'93 - '96	'90 - '04	Std. Err.
Equilibribum drift rate	μ_ξ	0.0116	-0.0554	0.0547	0.0401
Short-term mean-reversion rate	κ	1.5002	1.5624	1.2148	0.0270
Short-term risk premium	λ_χ	0.2740	0.0366	0.2758	0.0429
Short-term volatility	σ_χ	0.3411	0.1912	0.3614	0.0114
Equlibrium volatility	σ_ξ	0.1623	0.1026	0.1532	0.0059
Correlation in increments	$\rho_{\xi\chi}$	0.3519	0.1721	0.0427	0.0557
Equilibrium risk-neutral drift rate	μ_ξ^*	0.0100	0.0236	-0.0080	0.0027
Std. Dev. of error for Measurement Eq.	s_1	0.0408	0.0422	0.0271	0.0007
	s_2	0.0028	0.0086	0.0023	0.0007
	s_3	0.0042	0.0000	0.0080	0.0002
	s_4	0.0019	0.0027	0.0031	0.0005
	s_5	0.0051	0.0000	0.0218	0.0006
	s_6	n/a	0.0041	0.0183	0.0006

Table 5.3 – Comparison of Estimates by Timeframe

We also show a graphical comparison in Figure 5.4 of the forecast that would be generated in this latest case to that generated from the January 1990 – February 1995 fit. The first and most obvious conclusion is that the forecasts start from very different places. In all cases shown in this section, forecasts start from values indicated by the last set of the state variables from the fit to the relevant data set. We also note that the current price is significantly above the long-term equilibrium level, but that level appears to have shifted upward based on the values seen in the out years. All of the parameter estimates in the 1993 to 1996 period fall outside the 95% confidence interval around the estimates from this study, as do the short-term mean reversion rate, correlation, and equilibrium risk-neutral drift rate estimates from the 1990 to 1995 study, further signaling significant changes in the parameters when we consider the extended time horizon.

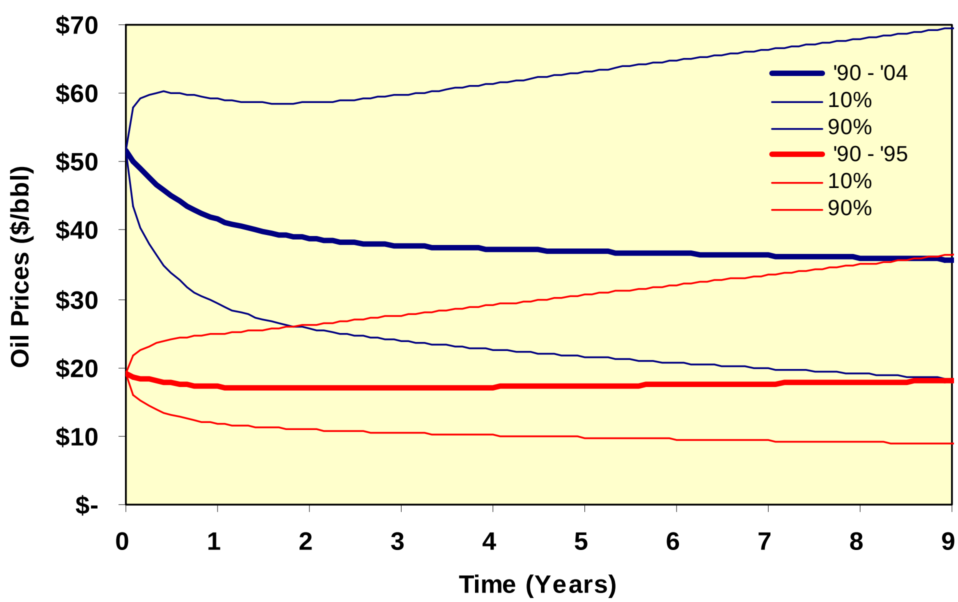

Figure 5.4 – Comparison of Forecasts Based on Different Time Horizons

Finally, to review the evolution of the underlying state variables ξ_t and χ_t, we show a plot of these two variables and the spot price given by them in Figure 5.5.

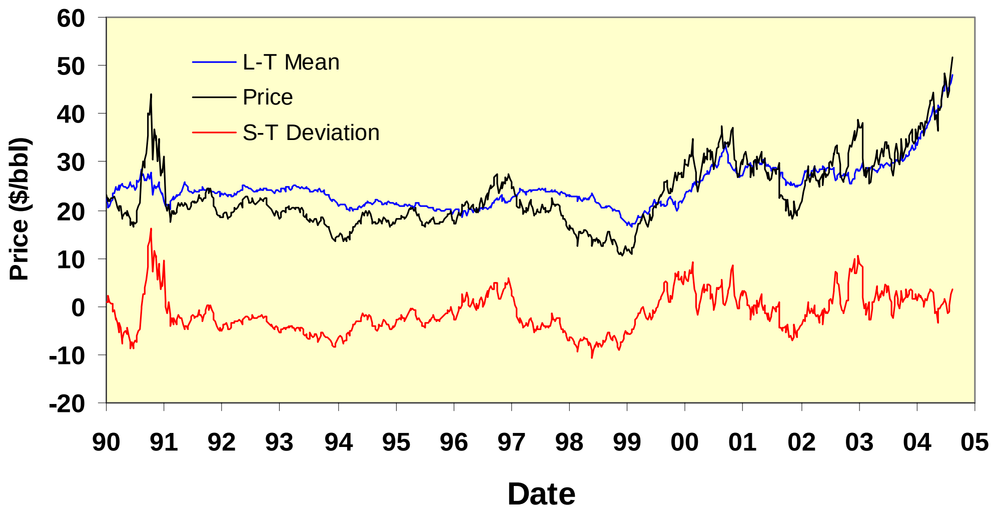

Figure 5.5 – Evolution of State Variables

All of the series in Figure 5.5 are unobservable; however as noted by Cox, Ingersoll, and Ross (1981), if the time to maturity for the futures price is relatively small, we can use it as a proxy for the spot price where it does not exist. We therefore show a plot in Figure 5.6 of the nearest futures maturity data with the spot price given by the underlying state variables to show that a good fit has been obtained. To evaluate the fit of an estimate to a data set, we can calculate the mean absolute percent error (MAPE). In this case, using our forecasted price as the estimate and the near-term futures price as the data set, we calculate a MAPE of 2.58%, indicating a good fit. We also note that the correlation between the two series is very high, having a value of 0.996.

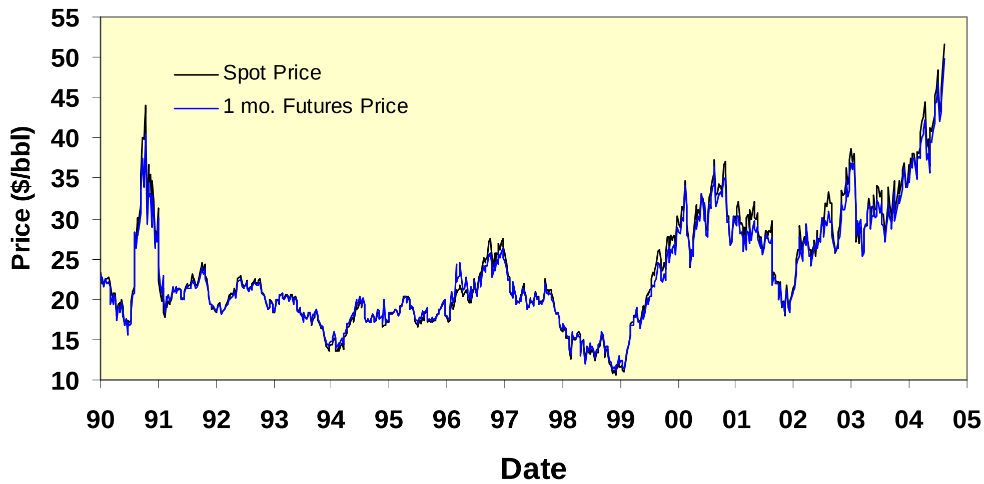

Figure 5.6 – Calculated Spot Price and Near-Term Futures Price

5.3 Estimation Using Seemingly Unrelated Regression

A second approach to determining the parameters for the Schwartz and Smith two-factor model is to use an alternative form of two-factor model, developed by Gibson and Schwartz (1990). Unlike the Schwartz and Smith model, this model is nested; the second factor is actually a parameter in the diffusion equation for the first factor. Therefore, it does not lend itself to straightforward modeling in a two-factor binomial lattice. However, estimation for this model is more straightforward, and uses nothing more than regression analysis. Schwartz and Smith (2000) showed that the two models are equivalent; therefore if we can find the parameters for the Gibson and Schwartz model, we can convert to the parameters needed for the Schwartz and Smith model.

The Gibson and Schwartz model includes a GBM process for X_t, the log of the spot price at time *t*, and an Ornstein-Uhlenbeck process for convenience yield, δ_t. Convenience yield is a parameter that measures the benefit to the holder of a commodity due to the option to sell it for consumption or use it in production. It fluctuates with supply and demand conditions and has been shown to be the primary factor in the relationship between spot and futures prices. It can be viewed as a dividend accruing to the commodity owner, as is the case in the Gibson and Schwartz process:

$$dX_t = (\mu - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1$$

$$d\delta_t = \kappa(\alpha - \delta_t)dt + \sigma_2 dz_2$$

where,

μ is the drift rate

δ_t is the convenience yield

σ_1 is the spot price volatiliy

κ is the mean reversion coefficient

α is the long-term convenience yield

σ_2 is the convenience yield volatility

Similar to the Schwartz and Smith model, the increments of the two processes are correlated:

$$dz_1 dz_2 = \rho dt$$

Seemingly unrelated regression (SUR) can be used to estimate ρ, as well as the parameters κ, α, and σ_2 for the convenience yield process. Convenience yields are calculated based on the relationship between the spot price and the futures price:

$$F(S_0, T) = X_0 e^{(r-\delta)\frac{T}{12}}$$

where T is time to futures maturity and r_T is the risk-free rate during the intervening time period from 0 to T, which leads to:

$$\delta = r_T - \frac{12}{T} \ln\left[\frac{F(S_0, T)}{X_0}\right]$$

For each period there is a spot price and corresponding futures prices of different maturities, so a time series for convenience yields can be constructed from data. Given the above mean-reverting process for the convenience yield, it can be shown that the expected value and variance of this distribution of δ are:

$$E[\delta_t] = \alpha + (\delta_0 - \alpha)e^{-\kappa T}$$

$$Var[\delta_t] = \frac{\sigma_2^2}{2\kappa}\left(1 - e^{-2\kappa T}\right)$$

These definitions and the discrete-time first-order autoregressive form of this Ornstein-Uhlenbeck process:

$$\delta_t - \delta_{t-1} = \underbrace{\alpha\left(1 - e^{-\kappa}\right)}_{a} + \underbrace{\left(e^{-\kappa} - 1\right)}_{b}\delta_{t-1} + e_t \quad \text{, or}$$

$$\delta_t - \delta_{t-1} = a + b\delta_{t-1} + e_t$$

can be used to formulate a regression to determine the parameters a and b. This regression is performed in conjunction with the regression for the price process

$$\ln\left(X_t / X_{t-1}\right) = a' + b' \ln\left(X_{t-1} / X_{t-2}\right) + \varepsilon_t$$

to capture the correlation between the two processes.

Using the same data set that was used in the previous section we ran a SUR model on both convenience yield and price using a common econometrics software package, LIMDEP. A separate regression with convenience yield was run for each futures maturity date, for a total of six different runs. For each run, the SUR model produces regression coefficients for each equation, as well as a 2×2 covariance matrix Σ for the residuals:

$$\begin{bmatrix} \sigma_{ee} & \sigma_{e\varepsilon} \\ \sigma_{\varepsilon e} & \sigma_{\varepsilon\varepsilon} \end{bmatrix}$$

The results from each run are summarized below in Table 5.4.

Maturity	**1 mo.**	**3 mo.**	**5 mo.**	**9 mo.**	**13 mo.**	**17 mo.**
Results from SUR:						
a	-0.9574	-0.2199	-0.0631	-0.0296	-0.0521	-0.0187
b	0.0587	0.0152	0.0069	0.0033	0.0056	0.0019
σ_e	0.1679	0.1421	0.0714	0.0453	0.0433	0.0285
$\sigma_{\varepsilon e}$	0.0010	0.0044	0.0027	0.0020	0.0015	0.0014
σ_ε	0.0525	0.0527	0.0529	0.0528	0.0525	0.0528

Table 5.4 – Regression Results for Different Maturities

With a and b determined from the convenience yield regression, κ, α, and σ_2 are calculated as follows:

$$k = -\ln(1 + b)$$

$$\alpha = -\frac{a}{b}$$

$$\sigma_2 = \sigma_e \sqrt{\frac{\ln(1+b)}{(1+b)^2 - 1}} .$$

The correlation ρ is calculated by using either off-diagonal term in covariance matrix Σ and dividing by the standard deviations of the residuals:

$$\rho = \frac{\sigma_{e\varepsilon}}{\sigma_e \sigma_\varepsilon}$$

The calculated parameters are shown for each maturity date in Table 5.5. The table shows significant variance in the parameter value, depending on the time to maturity.

Maturity	**1 mo.**	**3 mo.**	**5 mo.**	**9 mo.**	**13 mo.**	**17 mo.**
Parameters:						
α	0.0613	0.0691	0.1101	0.1110	0.1068	0.1022
κ	22.7510	1.7904	0.4699	0.2166	0.3862	0.1358
σ_2	0.2985	0.1132	0.0521	0.0325	0.0314	0.0204
ρ	0.1171	0.5829	0.7208	0.8479	0.6757	0.9077

Table 5.5 – Parameter Estimates for Different Maturities

Next, we determine the parameters σ_1 and μ directly from the futures price data. Depending on the frequency of data, the parameters can be adjusted to the desired reference period. If σ_T is the standard deviation for convenience yield of frequency T, to obtain σ on an annual basis the following conversion is used:

$$\sigma = \sigma_T \sqrt{T \, periods \ per \ year}$$

Volatility of the spot price, σ_1, can thus be determined directly from spot price data by tabulating the returns, $R_t = \ln(X_t / X_{t-1})$, and then computing the standard deviation. The drift rate, μ, can then be computed from the average of R_t , ν, and transforming to a lognormal mean:

$$\mu = \nu + \frac{\sigma_1^2}{2}.$$

Analysis of the data in our case showed σ_1 to be 38.4 % and μ to be 3.6 %

As was the case with the Schwarz and Smith model, the Gibson and Schwartz model can be used for valuation provided it is adjusted to reflect a risk-neutral forecast. For the Gibson and Schwartz model, this requires adjustment of the drift rate of the convenience yield process to account for the market price of convenience yield risk. The model specification then becomes:

$$dX_t = (\mu - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1^*$$

$$d\delta_t = [\kappa(\alpha - \delta_t) - \lambda]dt + \sigma_2 dz_2^*$$

$$dz_1^* dz_2^* = \rho dt$$

where λ is the market price of convenience yield risk. This parameter can be estimated, once the other parameters are determined, by using the futures valuation equation:

$$\ln(F_T) = X_T - \delta \frac{1 - e^{-\kappa T}}{\kappa} + A(T)$$

where,

$$A(T) = \left(r - \left(\alpha - \frac{\lambda}{\kappa} \right) + \frac{1}{2}\frac{\sigma_2^2}{\kappa^2} - \frac{\rho\sigma_1\sigma_2}{\kappa} \right) T + \frac{\sigma_2^2}{4}\frac{1 - e^{-2\kappa T}}{\kappa^3} + \left(\left(\alpha - \frac{\lambda}{\kappa} \right)\kappa + \rho\sigma_1\sigma_2 - \frac{\sigma_2^2}{\kappa} \right) \frac{1 - e^{-\kappa T}}{\kappa^2}$$

and solving for the value of λ that yields the best fit with the current futures price data. This was done in our case by setting up a simple Excel worksheet to calculate the futures prices and then using the add-in function Solver to find the value of λ that minimizes the squared deviations from the actual observed futures prices for each date. Using this approach, the values for λ ranged from -0.11 to 0.39, depending on the slope of the individual futures curve, however the average value was 0.185.

With estimates for all the parameters for the Gibson and Schwartz model we can use the following mapping to obtain the parameters for an equivalent Schwartz and Smith model:

Schwartz & Smith Parameter $\Rightarrow$ Calculated from Gibson & Schwartz Parameters

$$\kappa \Rightarrow \kappa$$

$$\sigma_{\chi} \Rightarrow \frac{\sigma_2}{\kappa}$$

$$\mu_{\xi} \Rightarrow \mu - \alpha - \frac{1}{2}\sigma_1^2$$

$$\sigma_{\xi} \Rightarrow \sqrt{\sigma_1^2 + \frac{\sigma_2^2}{\kappa^2} - \frac{2\rho\sigma_1\sigma_2}{\kappa}}$$

$$\rho_{\xi\chi} \Rightarrow \frac{\rho\sigma_1 - \frac{\sigma_2}{\kappa}}{\sqrt{\sigma_1^2 + \frac{\sigma_2^2}{\kappa^2} - \frac{2\rho\sigma_1\sigma_2}{\kappa}}}$$

$$\lambda_{\chi} \Rightarrow \frac{\lambda}{\kappa}$$

$$\lambda_{\xi} \Rightarrow \mu - r - \frac{\lambda}{\kappa}$$

The resulting parameters for the Schwartz and Smith model, shown in Table 5.6, indicate that the estimates vary considerably with the futures maturity.

Maturity	1 mo.	3 mo.	5 mo.	9 mo.	13 mo.	17 mo.
Calculated Schwartz & Smith Parameters:						
κ =	22.7510	1.7904	0.4699	0.2166	0.3862	0.1358
σ_χ =	0.0131	0.0632	0.1110	0.1500	0.0814	0.1500
μ_ξ =	-0.0083	-0.0161	-0.0571	-0.0580	-0.0538	-0.0491
σ_ξ =	0.3824	0.3507	0.3134	0.2686	0.3342	0.2555
$\rho_{\xi\chi}$ =	-0.0832	-0.4576	-0.5285	-0.6528	-0.5324	-0.7765
λ_χ =	0.0081	0.1033	0.3937	0.8539	0.4791	1.3619
λ_ξ =	0.0685	-0.0267	-0.3170	-0.7773	-0.4024	-1.2852

Table 5.6 – Parameter Estimates from Mapping

This raises the issue of deciding which futures maturities to use with this method. As noted by Schwartz and Cortazar (1994), stochastic process movements have an impact on futures returns across all maturities and are important in explaining return variance. Therefore it is important in estimating the stochastic process of prices to use a wide spread of information across all futures maturities. However, it is not clear how to accomplish that in this case, and a simple average of the parameter estimates from the different maturities yields significantly different answers than those from the previous section. We can select the estimates based on the three month maturity prices as the most similar set of parameters to those determined from Kalman filter estimation and plot forecasts for comparison, as shown in Figure 5.7.

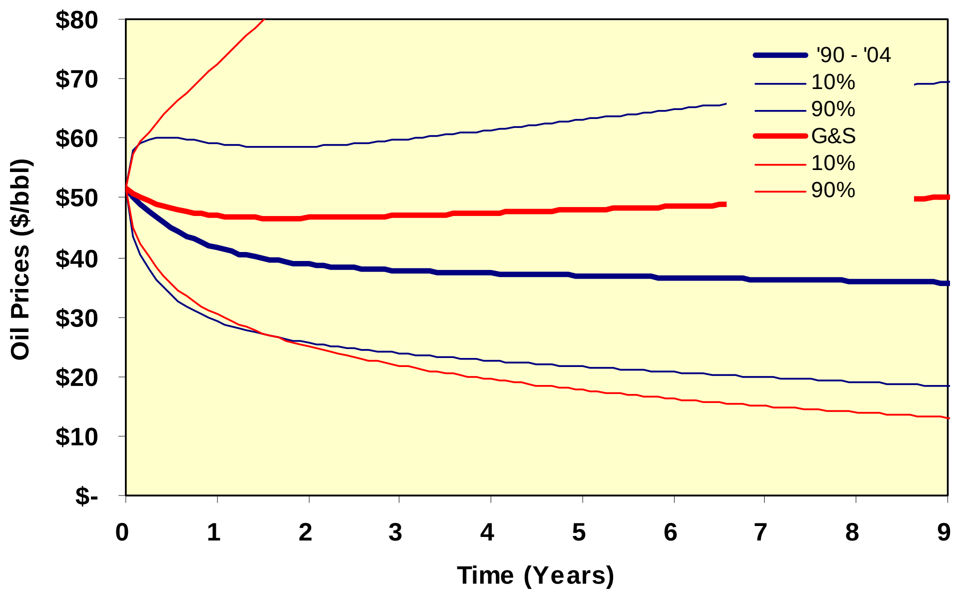

Figure 5.7 – Forecasts from Kalman Filter Estimation and Mapping

This figure shows that the two forecasts are still significantly different, both in terms of drift and variance. Furthermore, all of the parameter estimates except equilibrium drift rate are outside the 95% confidence intervals for estimates from the Kalman Filter estimates for this period. Therefore, although the regression approach with mapping to the Schwartz and Smith model parameters is simpler computationally relative to Kalman Filter estimation, in this case it did not provide a consistent set of parameters. This is likely due to the fact that this approach does not simultaneously consider futures data across the different maturities, as in the Kalman Filter approach. We therefore recommend that this method only be used when approximate estimates are required.

5.3 Implied Estimation Using The Current Futures Strip

The third method for estimating parameters for the Schwartz and Smith two-factor model is to fit the current futures curve with the futures pricing equation:

$$\ln(F_{T,0}) = e^{-\kappa T}\chi_0 + \xi_0 + A(T)$$

where,

$$A(T) = \mu_\xi^* T - (1 - e^{-\kappa T})\frac{\lambda_\chi}{\kappa} + \frac{1}{2}\left((1 - e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2(1 - e^{-\kappa T})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \right).$$

Using this equation and initial estimates for each of the parameters, the state variables that minimize the squared deviations between the calculated and observed futures prices can be obtained, thereby yielding the spot price. Based on the fit to the observed prices, the parameters can be changed and the process repeated until the fit cannot be further improved, in which case the final set of parameters become the estimates. This process is obviously easier to implement when some of the seven parameters are known, and only two or three of the seven parameters need to be estimated. Two parameters that can be estimated beforehand from historical data are the short-term and equilibrium volatilities. Schwartz and Smith (2000) propose using this method as a shorthand way to approximate the solutions obtained through the full Kalman Filter method

In this case, a simple worksheet was set up to implement the above pricing equation and compile the squared differences from the actual futures curve. Given a set of estimated parameters, Excel Solver was then used to find the state variables that minimize the sum of the squared differences. The objective function and constraints for this problem can be written as:

$$\min \sum \left(F_T - f_T\right)^2 \qquad \forall T$$

$$s.t. \;\; \xi + \chi = \text{log(current price)}$$

where:

F_T = calculated futures price using the equation above

f_T = observed futures price

T = futures maturity

A snapshot of the worksheet for this process is shown in Figure 5.8.

Stochastic Process Parameters

μ_ξ =	0.050
κ =	1.500
λ_χ =	0.200
σ_χ =	0.350
σ_ξ =	0.150
$\rho_{\xi\chi}$ =	0.300
μ_ξ^* =	-0.010
Δt =	0.083

State Values

ξ_0	χ_0
3.82403	0.08318

L-T Mean	Spot Price
45.788	49.760

		Futures Prices					
Date	**Spot**	30	90	150	270	390	510
9/28/04	49.76	49.90	48.49	47.17	44.61	42.65	41.02

Maturity	$\varepsilon^{-\kappa T}\chi_0+\xi_0$	**A(T)**	**ln($F_{T,0}$)**	F_T	**Obs.** F_T	Δ	
1.000	3.897	-0.010	3.888	48.795	49.900	1.221	
3.000	3.881	-0.027	3.854	47.175	48.490	1.728	
5.000	3.869	-0.042	3.827	45.904	47.170	1.602	
9.000	3.851	-0.064	3.787	44.136	44.610	0.224	
13.000	3.840	-0.078	3.763	43.066	42.650	0.173	
17.000	3.834	-0.086	3.748	42.423	41.020	1.970	
						6.919	minimize

Observed vs. Calculated Futures Price

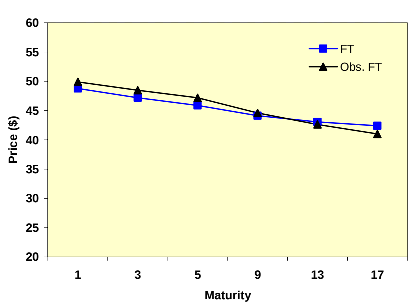

Figure 5.8 – Worksheet for Implied Method

We can select a set of parameters to estimate by running solver again with these state variables. For example, we can use historical estimates for σ_ξ and σ_χ, assume $\rho_{\chi\xi}$ is a known stable estimate, and then estimate the remaining parameters, κ, λ_χ, and μ_ξ^*. Figure 5.9 shows the worksheet with updated parameters and fit in this case, where we can see that the sum of squared errors has been substantially reduced from the previous case.

Stochastic Process Parameters

μ_ξ =	0.050
κ =	0.653
λ_χ =	0.161
σ_χ =	0.350
σ_ξ =	0.150
$\rho_{\xi\chi}$ =	0.300
μ_ξ^* =	-0.030
Δt =	0.083

State Values

ξ_0	χ_0
3.82403	0.08318
45.788	49.760

Date	Spot	Futures Prices 30	90	150	270	390	510
9/28/04	49.76	49.90	48.49	47.17	44.61	42.65	41.02

Maturity	$\varepsilon^{-\kappa T}\chi_0+\xi_0$	A(T)	$\ln(F_{T,0})$	F_T	Obs. F_T	Δ	
1.000	3.903	-0.009	3.894	49.122	49.900	0.606	
3.000	3.895	-0.025	3.870	47.921	48.490	0.324	
5.000	3.887	-0.041	3.846	46.818	47.170	0.124	
9.000	3.875	-0.071	3.804	44.884	44.610	0.075	
13.000	3.865	-0.098	3.767	43.268	42.650	0.383	
17.000	3.857	-0.121	3.736	41.919	41.020	0.808	
						2.319	minimize

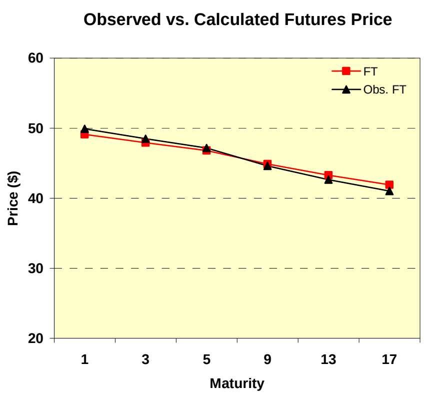

Figure 5.9 – Updated Worksheet for Implied Method

Figure 5.10 shows the forecast that would be generated from these parameters as compared to the parameters from the Kalman Filter approach.

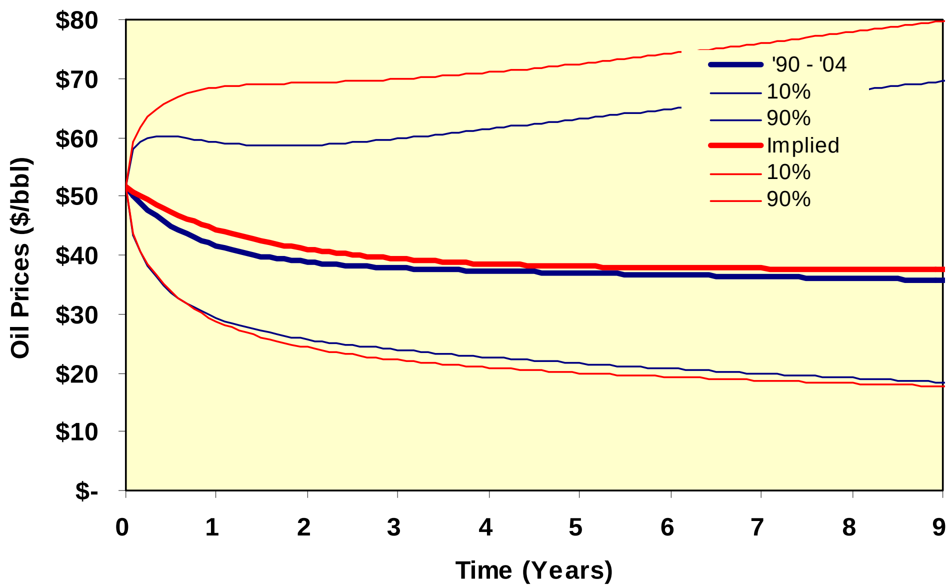

Figure 5.10 – Forecasts for Kalman Filter Estimation and the Implied Method

Although simple to implement, this process is rather ad-hoc in both the way parameter values are selected for estimation and in the way an "optimized" set of parameters is obtained. Furthermore, without having some knowledge of the parameter estimates beforehand, or having historical estimates of parameter values, it might be very difficult to find the optimum or near-optimum set.

5.4 SUMMARY

To summarize, in this section estimates for the two-factor Schwartz and Smith model were be obtained using three different methods. Since this evaluation was primarily undertaken to support testing of the models we have developed, we have not presented a detailed statistical comparison of the different approaches. However, we have documented our experiments with each approach. The Kalman Filter method provides not only parameter estimates, but also errors for those estimates, and is generally considered the best approach for this type of problem. However, it also requires considerable knowledge of its algorithm and its computational implementation, as well as facility with an appropriate platform in which to carry out the calculations. The other two approaches, especially the implied approach, impose a much lighter computational burden and use analysis tools familiar to most practitioners. The approach using a mapping from parameters determined through regression for the Gibson and Schwartz model did not provide results that were consistent with those from Kalman filtering, however. As a result, if an approach with lighter computational burden is required, the implied approach would be preferred. The implied approach also provides the benefit of a forward-looking analysis, and may be preferable when forecasting in the near term is the objective. Like the Kalman Filter, the implied approach can also return the state variables so that their evolution over time can be checked against near-term futures data.

6. APPLICATION TO A SWITCHING OPTION

The previous sections have detailed how the bivariate binomial approximation can be used to model a two-factor mean reverting processes for a single underlying asset. In this section the approximation is applied to a real option problem that has two underlying assets that both follow mean-reverting stochastic processes. Assuming both processes can be modeled as one-factor Ornstein-Uhlenbeck processes, the bivariate approximation applies as outlined at the end of Section 3.

6.1 DESCRIPTION OF THE APPLICATION

The application deals with valuing an important research and development prospect in the area of enhanced oil recovery. This project has implications for development scenarios for one of North America's largest producing areas, the North Slope of Alaska, which currently comprises 25% of the total U.S. oil production and 30% of its remaining oil reserves. The technology to be evaluated is low-salinity water flooding, which has the potential to increase the amount of oil recovered by up to 10% relative to conventional waterflooding techniques, as discussed by Webb et. al. (2004) and McGuire et. al. (2005). Conventional waterflooding is the practice by which water is injected into an oil reservoir via dedicated injector wells to artificially maintain the reservoir drive mechanism of water sweeping oil toward producer wells. Typical recovery percentages under waterflooding in the North Slope can reach nearly 60% of the original oil in place, leaving approximately 40% of the oil behind. Given that the oil in place was 55 billion barrels (Bbo), an incremental 10% in ultimate recovery would have a huge economic impact. However, the North Slope also holds an estimated 35 trillion cubic feet (Tcf) of natural gas, some of which is currently being produced with oil and re-injected into the reservoir because there is no pipeline to transport it to North American markets. A pipeline which would cost an estimated $19 billion and would follow a route through northwestern Canada is currently being evaluated, both in terms of its economic viability and its technical and regulatory feasibility. If gas production commences, the assumption is that oil production would be impacted. The current producing wells are optimally configured and operated to maximize oil production; therefore a reconfiguration or change in operational approach would reduce the oil production at the expense of increasing gas production. Furthermore, energy in the form of gas pressure is removed from the reservoir rather than replaced, as is the current practice.

Given this context, a real option analysis can be carried out to determine the value of a research and development project to evaluate low-salinity waterflooding. This value will naturally be contingent upon the optimal course of action for managing North Slope production and timing the Alaska gas pipeline. The optimal course of action will be determined by the relationship between oil and gas prices, as well as the decline of continued oil production from a finite-sized reservoir. The point of switching from only oil production to combined oil and gas production could thus be affected by the success of low-salinity waterflooding in stemming the oil production decline. At the same time, the economic viability of pursuing low-salinity waterflooding depends on the remaining length of time, and thus volume, of oil production before the switch to gas production. Therefore there is a classic recursive relationship between future production and the low-salinity project that requires optimization through dynamic programming techniques.

6.2 RELATED RESEARCH

The background literature on the issues involved in this problem is varied. Several different types of options, such as exchange options and rainbow options, have similar characteristics to this problem, although there are also key differences in each case. Copeland and Antikarov (2001) present a simple example of a switching option between two modes of operation for a factory, and they obtain a solution using a discrete lattice. They assume two GBM process for the uncertain cash flows from the two modes of operation, and approximate with the standard binomial lattice, assuming that the two processes are uncorrelated. The authors also note that both a correct valuation of the project with flexibility and the optimal management policy are obtained through their analysis. Cases where the underlying asset is exhaustible are suggested as an area for extension for this type of approach. Bailey et. al. (2003) discuss several applications of real options, including a switching option for the size of a key processing component during the design phase of a facility construction project. They propose using a similar discrete time approach to that used by Copeland and Antikarov (2001), with independent lattices for the two facility size options.

There are several studies on the optimal extraction of a depletable natural resource base, including Brennan and Schwartz (1985), Dixit and Pindyck (1994), Carlson, Khokher, and Titman (2000), Dias, Rocha, and Teixeira (2003), and Ronn (2004) that are useful for developing analytical representations of this problem, however they do not address the case of more than one coexisting resource.

Outside of the natural resources literature, Margrabe (1978) developed an early model for valuing the option to exchange one asset for another; however this closed-form solution is valid only for the case of underlying assets that follow GBM diffusions and also does not consider depletion. Carr (1988) generalizes this model to the case of sequential exchange opportunities. Dixit (1989) studies entry and exit decisions of a firm under uncertainty, developing closed-form solutions for the value with this flexibility under a GBM assumption and also considering the case of an underlying one-factor mean-reverting process. In the latter case Dixit notes that closed form solutions cannot be obtained.

Other related research on problems with similar characteristics includes the work of Stulz (1982), Johnson (1987), and Boyle and Tse (1990) who find analytical solutions for pricing options on the maximum of multiple assets, but only under the assumptions of underlying GBM diffusions and non-depletable assets. Childs, Ott, and Triantis (1998) also investigate valuation for multiple assets and specifically consider the case where the assets are interrelated. They develop a closed-form solution for the case of a European option, but note that a more realistic formulation to allow early exercise would require numerical approximation. Wilmott, Howison, and Dewynne (1995) assume a two-factor correlated process to develop a framework for valuing a convertible bond with stochastic interest rate, however it does not accommodate a depletable asset and in any case must be solved numerically.

6.3 DEVELOPMENT OF AN ANALYTICAL MODEL

We are not aware of an existing analytical solution that can be easily adapted to the problem being considered here. To evaluate whether a closed-form solution can be found for this problem, the general approach used by Dixit and Pindyck (1994) can first be used to set up the differential equation to model this problem. In this case, there will be four state variables for the problem:

O = price of oil

G = price of natural gas

R_O = oil reserves

R_G = gas reserves.

The stochastic process for oil price O and natural gas price G are assumed to be single-factor mean-reverting processes:

$$dO_t = \kappa_O(\overline{O} - O_t)]dt + \sigma_O dz_t^O$$

$$dG_t = \kappa_G(\overline{G} - G_t)]dt + \sigma_G dz_t^G,$$

with correlated process increments, $dz_t^O dz_t^G = \rho dt$, and reserves are expected to change over time according to the following relationships, which depend on which types of production are active:

$$\begin{cases} dR_O = -\alpha_O R_O dt \\ \quad dR_G = 0 \end{cases} \qquad \text{if only oil production}$$

$$\begin{cases} dR_O = -\alpha_O R_O dt \\ dR_G = -\alpha_G R_G dt \end{cases} \qquad \text{if both oil and gas production}$$

$$R_O(\tau - \delta t) = \omega R_O(\tau + \delta t) \qquad \text{at time of switching, } \tau$$

where,

κ_O, κ_G = coefficients for the speed of mean reversion

$\overline{O}, \overline{G}$ = long-term mean commodity prices

σ_O, σ_G = process volatilities

dz_O, dz_G = random increments of the processes

α_O, α_G = production decline rates, and

ω = factor for impact of gas production.

With these specifications and using the risk-free discount rate r, the unit value V' of the project during oil production only must satisfy the differential equation:

$$\alpha_O R_O \theta_O V' + \kappa_O \left(\overline{O} - O_t\right) \frac{\partial V'}{\partial O} + \frac{1}{2}\sigma_O^2 \frac{\partial^2 V'}{\partial O^2} - rV' = 0.$$

When oil and gas production both occur, the unit value V'' must satisfy:

$$\alpha_G R_G \theta_G V'' + \alpha_O R_O \theta_O V'' + \kappa_G \left(\overline{G} - G_t\right) \frac{\partial V''}{\partial G} + + \kappa_O \left(\overline{O} - O_t\right) \frac{\partial V''}{\partial O} +$$

$$\frac{1}{2}\sigma_G^2 \frac{\partial^2 V''}{\partial G^2} + \frac{1}{2}\sigma_O^2 \frac{\partial^2 V''}{\partial O^2} + \rho \sigma_G \sigma_O \frac{\partial^2 V''}{\partial O \partial G} - rV'' = 0$$

Under these conditions, we satisfy the Bellman equation for the optimal control policy for the project. Next, the particular characteristics of this problem need to be incorporated by specifying boundary conditions for the above differential equation. These are given by the following:

$R_O \geq 0$

$R_G \geq 0$

$V'(O, R_O) \geq V''(O, G, R_O, R_G) - K$, and

$V''(O, G, 0, 0) = 0$,

which set the economic stopping and strike conditions. In these equations, θ_O and θ_G are oil and gas production conversion factors, which use an assumed point-forward production decline profile to convert reserves to an unit factor that can be multiplied by price to determine present value.

This formulation is somewhat similar to that for a power plant input fuel switching option provided in Dixit and Pindyck (1994); however it is slightly more complicated due to the depleting oil reserve base, rather than finite time horizon, and the underlying mean-reverting processes. In that example they note that the resultant partial differential equation must be solved numerically and propose using the binomial method, a direction which is followed for the remainder of this section.

6.4 NUMERICAL SOLUTION

As a numerical approach for solving options with underlying mean-reverting processes, discrete trees have found some limited use. Slade (2001) uses binomial trees based on the Nelson and Ramaswamy (1990) approach to model a one-factor mean-reverting process for metals price in valuing options for a mining operation. In the area of financial options, Hull's (1994) trees are used extensively for valuing interest rate derivatives and Jaillet, Ronn, and Tompaidis (2004) and Lari-Lavassani, Simchi, and Ware (2001) use binomial or trinomial trees to value swing options based on mean-reverting commodity prices. In the example problem being considered here, the binomial approximation method developed in Section 3.3 will be used.

6.4.1 Definitions and Assumptions

To solve the problem numerically, we need to first further define the problem and state the assumptions that will be used. First, although several firms share ownership in North Slope production, the project will be valued from the perspective of a single firm with one non-operating partner, the State of Alaska, which holds the mineral rights and therefore collects a 12.5% royalty. As a non-operating partner the State has 0% working interest and does not share in any costs. As is usually the case with real options problems, there are some required assumptions with regard to timing. The operating firm could in principle decide to switch production from oil to gas at any point in continuous time; however we will assume in this problem that this continuum is discretized into annual periods. The actual decision-making frequency in a firm is likely to be somewhere in between the two extremes. We will also assume that when a switch is made, it occurs instantaneously.

After a switch to gas production, we make the base-case assumption that the oil production would be reduced by 10%. This can be included in the model as a downstream private, or non-hedgeable, risk as an extension. Reserves for both for oil and natural gas are also assumed to be deterministic, but could also be modeled as private uncertainties in extended models. However, since oil has been produced from the North Slope since the early 1970's, there is a high level of certainty about reserves levels.

The low salinity technology has been tested in single wells with the impact measured by tracking chemical tracers injected with the water that was later recovered in nearby producing wells (McGuire et. al., 2005). Based on these pilot tests, the estimated probability that a more extensive test using a three-well grid will verify that an incremental 9% of oil can be recovered is 45%, while there is a 55% chance that larger scale testing will show that incremental recovery will only be 4% with low salinity water. These estimates are based on information from industry personnel familiar with the characteristics this technology. If the large scale test of sweep efficiency is successful, the next problem would be to prove that an operating-scale desalinization plant capable of producing the volumes of water needed could be feasible. The test plant will produce 50,000 barrels of water per day, and the estimated probability that it will operate efficiently is 40%, leaving a 60% chance that the plant will be inefficient. If the plant is inefficient, oil production would be impaired by 20% due to the lower volumes of water available for flooding the reservoir. The tests will be run in conjunction, since the desalinization plant is needed for the test waterflood.

6.4.2 Analytical Framework and Base Case Analysis

A base case solution to the problem can be obtained by finding the deterministic net present value using a simple decision tree or spreadsheet model with expected values for all inputs. In each year of the project the firm will decide whether to pay the switching cost to activate the pipeline and switch to gas production with reduced oil production, or to continue with producing oil and wait until the next year to revisit the switching question. Later in this section, we will remove some of the deterministic assumptions and compare to this base case.

Among the required inputs for this analysis are oil and gas prices and production profiles. The one-factor Ornstein-Uhlenbeck process is used in the example to model the separate processes. Although not as economically sophisticated as the two-factor models, Smith and McCardle (1999) use a one-factor mean-reverting process in their analyses of the application of option valuation techniques to oil and gas projects and note that it provided a fit to the empirical data in their case. The parameters for the two processes were determined using the same implied approach discussed in the previous chapter, and are summarized in Table 6.1.

	Oil		Gas	
Current Spot Price	53.00	$/bbl	6.50	$/Mcf
Mean-Reversion Coefficient, κ	0.5		2.0	
Long-Term Mean Price	45.00	$/bbl	4.50	$/Mcf
Process Variance, σ	20%		30%	

Table 6.1 – Process Parameters for Oil and Gas

Pindyck and Rotemberg (1990) discuss the correlation between commodity prices as a well known phenomenon, and this was also empirically observed by Moel and Tufano (1998) and others. In this case, the correlation between oil and gas prices was estimated to be approximately 30%, based on data from 1990 through 2004. Using the parameters from the above table yields the deterministic forecasts for oil and gas over a ten-year period shown in Figure 6.1. The figure shows that gas reverts more quickly to its long-term mean than oil does, stabilizing near its long-term equilibrium level by the third period. Both commodities are currently well above their estimated long-term means. Although the confidence intervals around the expected values are shown as in Section 5, only the expected value forecast lines, shown in bold in the figure, are relevant for the base case deterministic analysis.

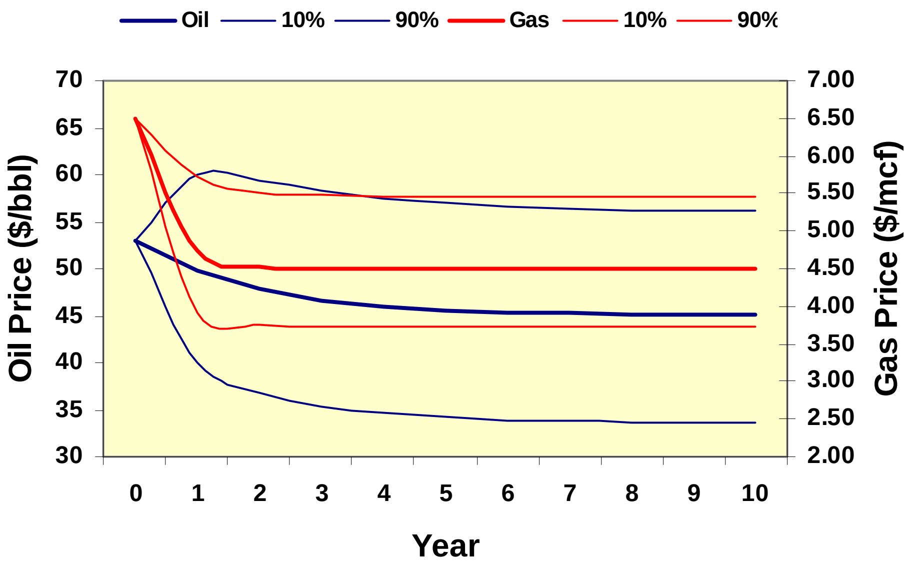

Figure 6.1 – One-Factor Price Forecasts for Oil and Gas

NPV calculations are greatly simplified with the use of a reserves factor that converts a given reserves amount to an assumed production profile. These factors can be used with the price forecast to calculate a series of cash flows which yield a present value. Implicit in the reserves factor are decline and discount rates. For this example, we assume that the historical decline rate for North Slope oil production holds, which is approximately 5% in an exponential model, so that the amount of reserves R remaining at any time t is given by:

$$R = R_0 e^{-0.05t} .$$

There is no historical gas production decline information for the North Slope, since the gas has been re-injected. However, standard reservoir engineering calculations based on the pressure, temperature, volume and rock permeability of the reservoir indicate an estimated exponential decline rate of 4.25%. The initial oil production rate is assumed to be approximately 1.6 million barrels per day, and initial gas production is estimated to be 4 billion cubic feet per day based on reservoir engineering calculations and pipeline capacity. These assumptions give the following production profiles that are shown in Figure 6.2.

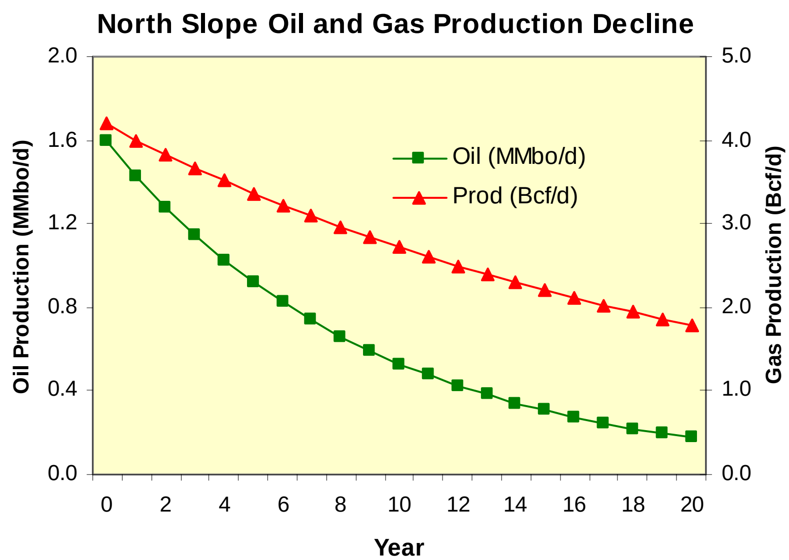

Figure 6.2 – Production Forecasts for Oil and Gas

We assume a 5% discount rate for deriving the reserves factors, since this is the approximate risk-free discount rate and the hypothetical firm should have a risk-neutral view of the private uncertainties that affect production forecasts. With these assumptions about decline and discounting, the reserves factors are obtained by forecasting a unit of production, assuming a price forecast, and then calculating the present value at the chosen discount rate. Using this approach, the reserves factors for oil and gas for this example were calculated to be $\theta_O = 0.6699$ and $\theta_G = 0.4248$, respectively. This means that, for example, the value of one barrel of oil reserves, or a "barrel of oil in the ground", when the current price is \$40 is \$26.80. Such factors are commonly used in rule-of-thumb estimation of oil and gas property values for screening acquisition and divestiture opportunities.

The net present value of the two alternatives at each step, continuing oil production or commencing oil/gas production, can then be calculated using:

$$PV_t = \max\left\{\left(P_O(t) - \Delta_O\right)Q_t(1-\lambda)\delta_s\delta_p + PV_{t+1}e^{-rt}, PV_{switch}\right\}$$

where,

$$PV_{switch} = \left\{\left(P_G(t) - \Delta_G\right)R_G(1-\lambda)\theta_G - K\right\} + \omega\left\{\left(P_O(t) - \Delta_O\right)R_O e^{-\alpha t}(1-\lambda)\theta_O\delta_s\delta_p\right\}$$

$P_O(t), P_G(t) =$ Prices of oil and gas at time t

$\Delta_O, \Delta_G =$ Price differentials due to processing and transportation costs

$Q_t =$ Oil Production rate

$\lambda =$ State of Alaska royalty

$\delta_s, \delta_p =$ Efficiency factors for low salinity sweep and desalinization plant

$R_O, R_G =$ Reserves for oil and gas

$K =$ Switching cost (cost of pipeline)

$\omega =$ Impact on oil production due to gas production

$\alpha =$ Exponential decline coefficient for oil production rate

This equation represents the optimization between the value of switching and the value of continuing the recursion for another step, and is easily implemented in spreadsheet or decision tree format. A simple decision tree constructed using DPL software is shown in Figure 6.3.

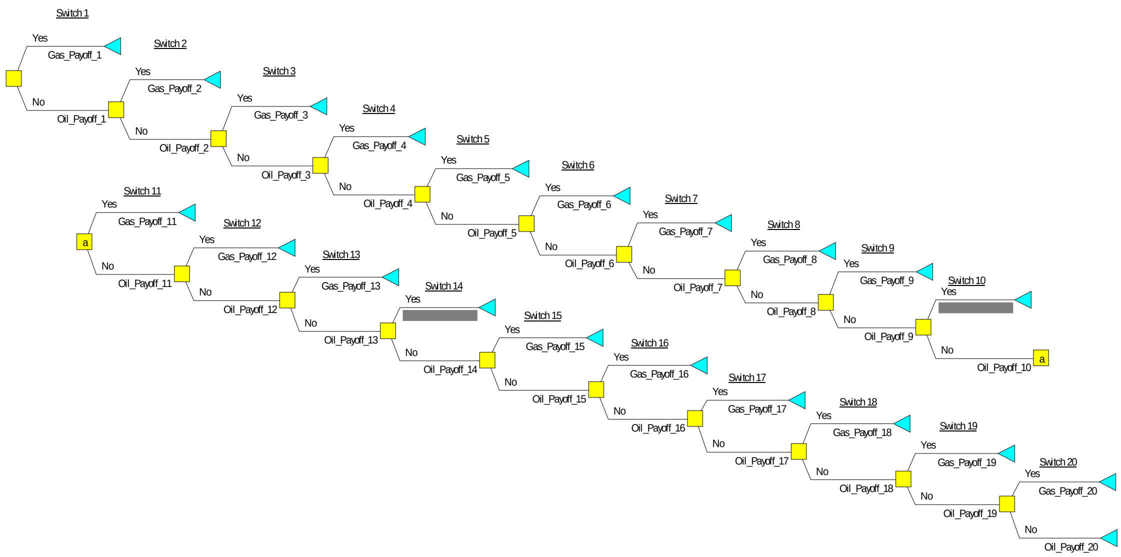

Figure 6.3 – Decision Tree for Switching Problem

Using the inputs shown in table 6.2 yields the policy summary shown in Figure 6.4.

Inputs	Oil		Gas	
Current Spot Price	53.00	$/bbl	6.50	$/Mcf
Transportation/Processing Cost, Δ	5.0	$/bbl	1.0	$/Mcf
Reserves	35.00	Tcf	5.00	Bbo
Oil decline rate, α	11%			
Oil Production Rate	1.43	MMbo/d		
Royalty, λ	0.125			
Gas production impact, ω	0.9			
Switching cost (pipeline)	19	$Bn		

Table 6.2 – Base Case Model Inputs

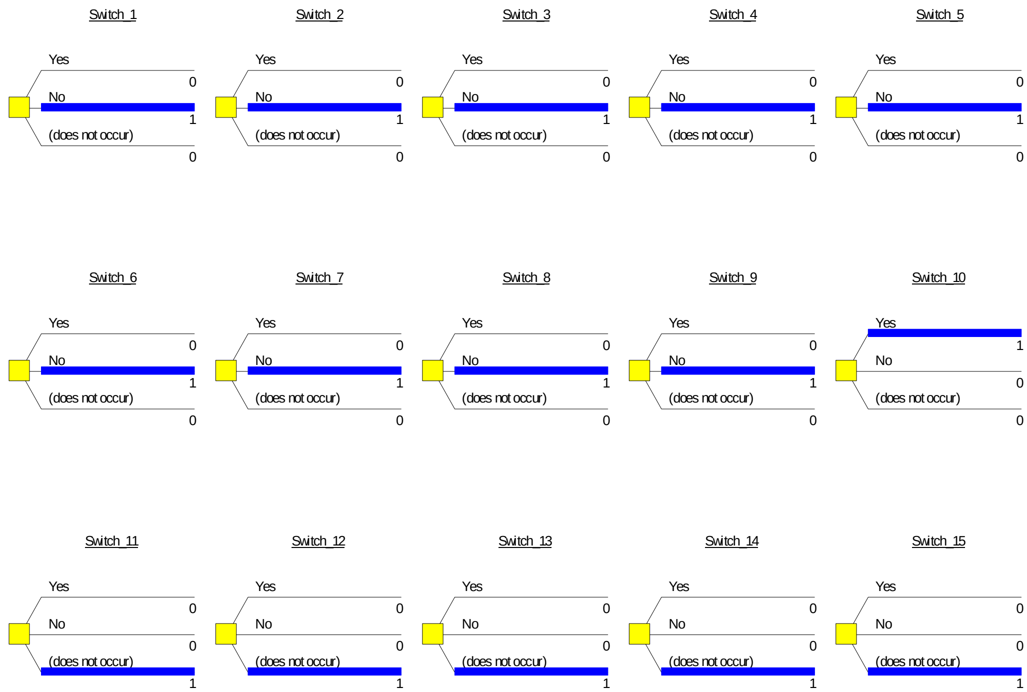

Figure 6.4 – Policy Summary for Deterministic Case

The figure indicates that the optimal point at which to switch to combined gas and oil production is year 10, which yields a Net Present value of $131.29 billion. In the deterministic success case, that is, when the low salinity process works with certainty and ultimate reserves recovery is 9% greater than the base case, the Net Present value increases to $141.84 billion.

6.4.3 Adding Technical Uncertainties

With a basic model of the problem constructed, we can now model the key underlying uncertainties. To assess which uncertainties should be modeled, we first consider the degree of uncertainty around selected variables, and then determine whether the solution will be sensitive to those uncertainties by performing sensitivity analysis with the base case solution. The variables and ranges selected for investigation in this case are oil price ($40 to $70), oil reserves (4.5 Bbo to 5.5 Bbo), oil decline rate (9% to 14%), gas price ($4 to $8), gas reserves (30 Tcf to 40 Tcf), gas production impact (0.8 to 1.0), low-salinity waterflood sweep efficiency gain (1.0 to 1.09), desalinization plant efficiency (0.8 to 1.0), and switching costs ($15 Bn to $25 Bn). The result of these sensitivities are shown in the Tornado diagram in Figure 6.5.

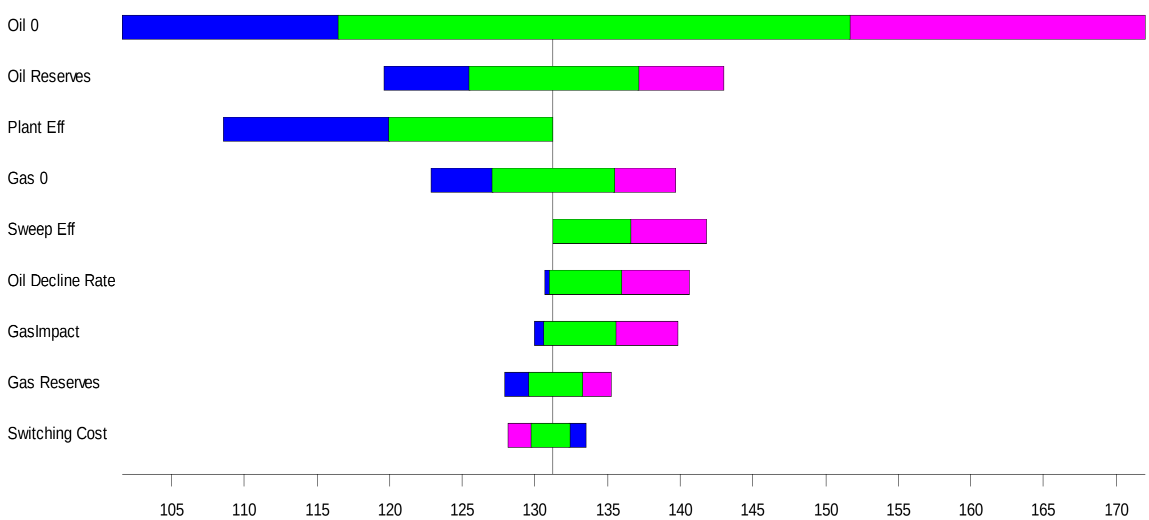

Figure 6.5 – Sensitivity Analysis for Base Case

This figure shows that oil and gas price as well as the two private technological uncertainties comprise four of the five most significant variables. The solution is also very sensitive to the oil reserves, even though there is a fairly high level of certainty around the estimate of this variable.

We thus begin construction of a dynamic model by incorporating the private uncertainties associated with low salinity waterflooding into the model. Following the outcome of the tests of plant and sweep efficiency, a decision will be made on whether or not to implement low-salinity waterflooding. This uncertainty-decision sequence is depicted in Figure 6.6.

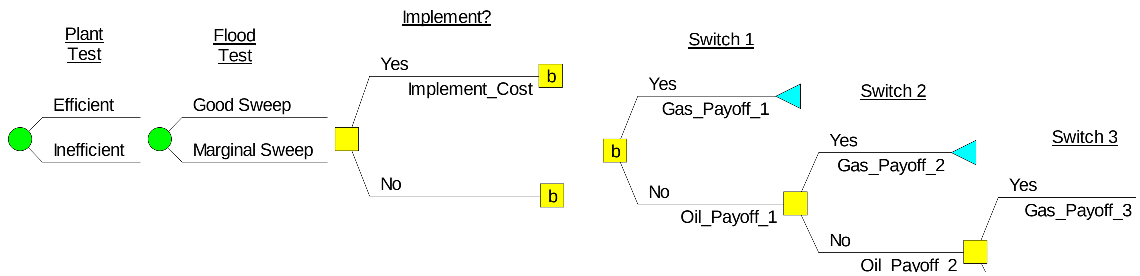

Figure 6.6 – Adding Technology Test Uncertainty

This sequence precedes the tree shown in Figure 6.3, denoted as sub-tree "b" in the above figure. The values for plant efficiency, δ_p and sweep efficiency, δ_s are contingent on both the test outcomes and the firm's implementation decision, as shown in Figure 6.7 below.

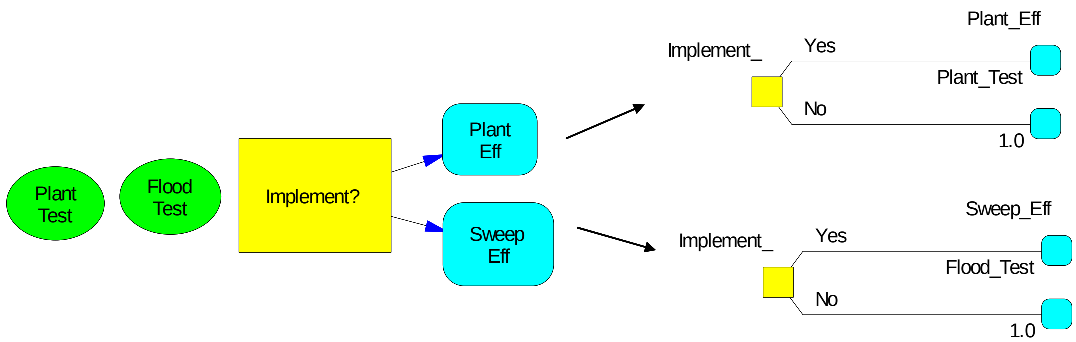

Figure 6.7 – Conditioning Efficiency Terms on Implementation

In cases where the firm decides not to implement, the efficiency terms are therefore equal to one, which is the benchmark conventional waterflood. The four cases and their expected outcomes are shown in Figure 6.8. When the low-salinity approach is implemented, an additional development cost of $5.50/bbl, or $2.5Bn is incurred on a present value basis over the life of the project. Therefore this cost is entered in the decision nodes where it applies.

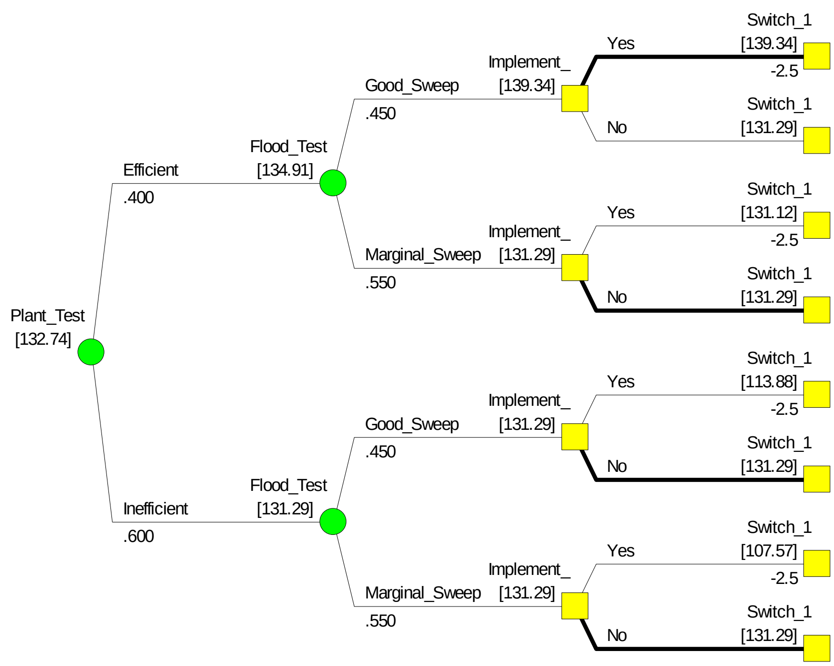

Figure 6.8 – Private Uncertainty Outcomes

As shown in Figure 6.8, the expected value is $1.45 Bn more than the base case, due to one outcome where the optimal decision is to implement, which produces an additional $8 Bn over the base case. The outcome of an efficient plant test coupled with a marginal improvement in sweep efficiency is nearly breakeven. In the case where low salinity waterflooding is implemented, the impact on the switch to gas production and the decision on when to construct the pipeline is to move it back from year 10 to year 13. This makes sense intuitively, since the oil production is extended further out in time with the added recoverable reserves.

The model in its current state is still deterministic with respect to commodity prices. The next step is then to add price uncertainty in the form of mean-reverting stochastic price processes for the two commodities.

6.4.4 Adding Commodity Price Uncertainties

Commodity price uncertainty can be added to the model by implementing the approach developed in Section 3 and using the parameters given in Table 6.1. Since low-salinity waterflooding has the primary objective of increasing oil production, the stochastic process for oil is added first. This is done by changing the deterministic value entries for oil price in each period to chance nodes with the probabilities and up/down movements specified per the one-factor binomial model developed in Section 3.1. These chance nodes are then added to the tree as shown in Figure 6.9 for the first three periods.

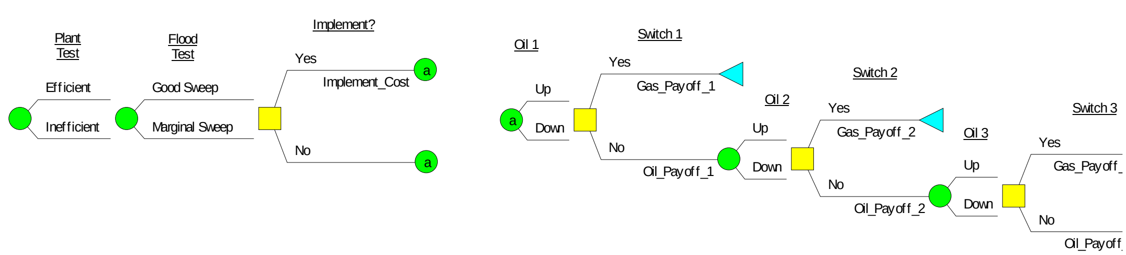

Figure 6.9 – Adding Oil Price Uncertainty

Solving this tree yields a value of $139.56 Bn which reflects increased value due to optimized decision-making under uncertainty, but also includes error due to the binomial approximation, as shown in Section 4.

The tree with the oil price uncertainty starts to become very large, with over one million endpoints for the 20 price nodes alone. With a view toward the additional nodes that will be needed to incorporate the gas price uncertainty, the tree can be economized by noting that in late periods, specifically in periods 15-20, the price has largely reverted to its equilibrium level. By this point, decisions have been made for most paths, and discounting minimizes the impact of any changes in cash flows as well. Modifying the tree to include constant value nodes rather than chance nodes in periods 15-20 reduces the number of endpoints by nearly 97% which drastically reduces computing time and still yields a value, $139.49 Bn, which is very close to the solution with the full tree.

The policy summary from this solution, shown in Figure 6.10, indicates a marked change in clarity about when the switch to oil and gas production should commence, due to the added uncertainty. Of note, the up move in the first period triggers an immediate move to oil and gas production. Although this seems counterintuitive, since oil price occurs in both the value for continuing only oil production and for commencing both oil and gas production, the impact of reduced oil production rate on value in the second scenario is more than offset by the increase in oil price. The effect of price uncertainty is complex, however, For example, there are also some paths where high prices occur later in time and offset the oil production decline to delay the switch to oil and gas. These cases can be noted in Figure 6.10, as there are some decisions to switch that do not occur until year 15.

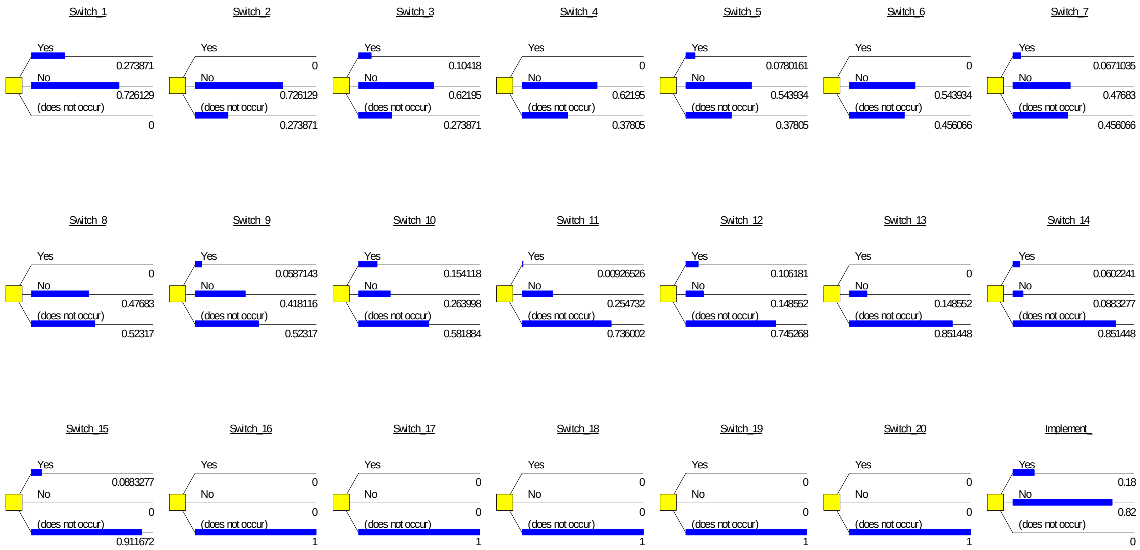

Figure 6.10 – Oil Price Uncertainty Effect on Policy

The last step to complete the model under the framework and assumptions used to this point is to add the gas price uncertainty. Following the same procedure as with oil price, the gas price uncertainty is added with a binomial approximation of a mean-reverting stochastic price process with the parameters given in Table 6.1. Based on what was learned from the oil price uncertainty modeling, and noting that the speed of mean-reversion for gas is much higher than for oil, the number of periods to be modeled with chance nodes can be reduced for gas as well. A review of the deterministic gas price forecast indicates that with the given parameters, it largely reverts to the long term equilibrium level within the first three years. Therefore, only the first three periods in the tree will be modeled with chance nodes. These chance nodes are then added to the tree as shown in Figure 6.11 for the first three periods.

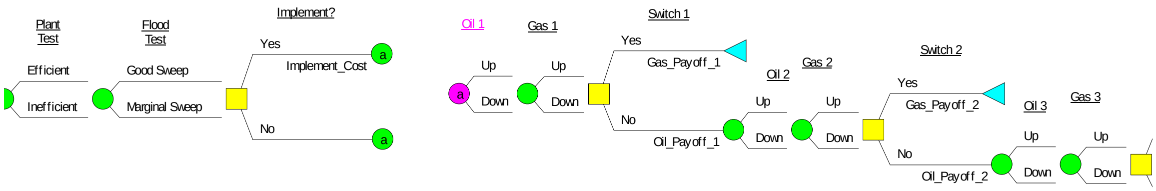

Figure 6.11 – Adding Gas Price Uncertainty

Solving this tree increases the value to $142.36 Bn, as shown in the partially expanded tree in Figure 6.12. This again reflects value due to optimized decision-making under the added uncertainty, as well as the error due to the binomial approximation.

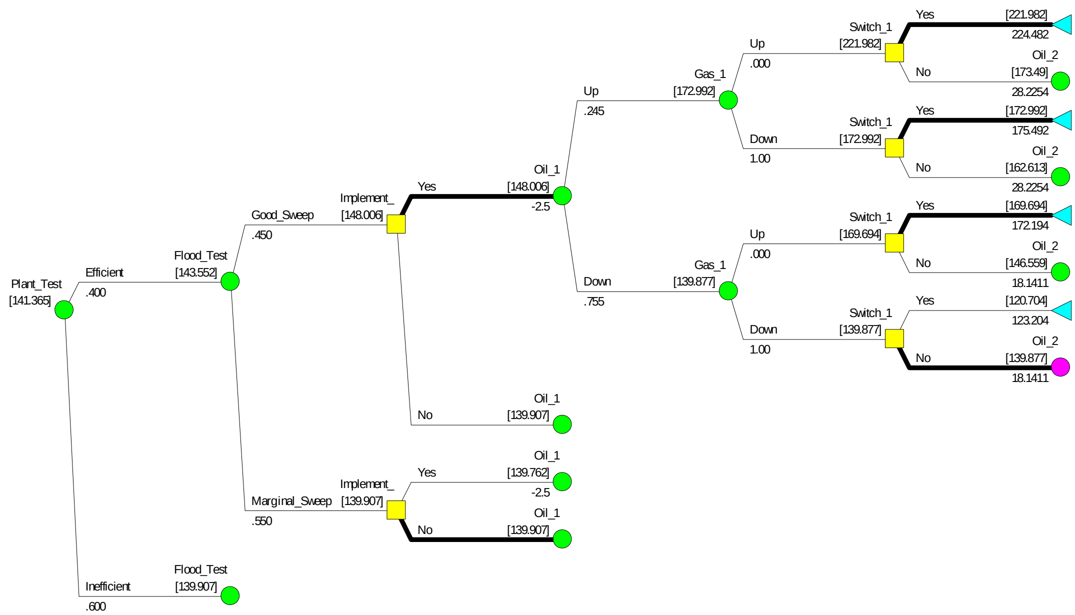

Figure 6.12 – Partially-Expanded Solved Decision Tree

As can be seen by comparing in Figure 6.13 to Figure 6.10, there are only subtle changes to the decision policy due to the addition of the gas price uncertainty. The first period exercise rate goes down only slightly, from 0.273 to 0.245, and the last period in which an exercise takes place is still period 15, although the rate goes up from 0.088 to 0.096.

Figure 6.13 – Policy Under Gas and Oil Price Uncertainty

We note at this point that there is a peculiar pattern in both Figures 6.10 and 6.13 where there are some periods with zero instances of an exercised decision. This is due to the typical convergence pattern seen in binomial approximations, as was discussed in Section 3.

6.4.5 Correlated Uncertainties

The two-factor model developed in Section 3.3 provides the capability to incorporate correlation between the two commodities. Correlation affects the distribution of outcomes in any increment through the probability calculations, and thus we expect it to have an effect on option values. By including the estimated correlation of 0.30 in the decision tree model, there is indeed a slight impact on the value, as it increases from \$141.37 Bn to \$142.36 Bn. For some insight as to the changes under correlation a plot of cumulative fraction of exercise is provided in Figure 6.14.

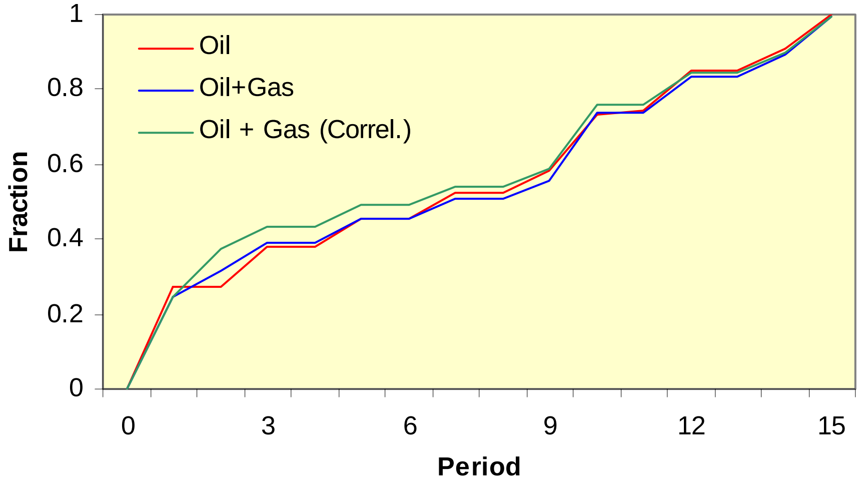

Figure 6.14 – Policy Summary: Three Incremental Cases

The figure shows that, although the curves are fairly similar, there are differences in the frequencies in periods 2 – 6. In particular, we observe that in the case of correlated uncertainties, the option to switch is executed earlier. From reviewing Figure 6.1, it follows that if the two commodity prices are moving in step and with their respective forecasts, then periods 1-6 is the period during which the rate of divergence is greatest and switching is likely to be triggered.

6.4.6 Summary and Interpretation of Results

To summarize the results to this point, the values calculated and model assumptions for each of the decision tree models are presented in Table 6.3.

Case	PV ($Bn)
Base (continued North Slope production without new technology, deterministic price forecast)	**131.29**
Success (100% chance of success for technological risks, deterministic price forecast)	**141.84**
Expected Value (expected value for technological risks, deterministic price forecast)	**132.74**
Oil price uncertainty (stochastic price forecast for oil)	**139.56**
Oil price uncertainty (stochastic forecast to 15 years, deterministic thereafter)	**139.49**
Oil and gas price uncertainty (stochastic price forecasts for both oil and gas)	**141.37**
Oil and gas price uncertainty with correlation (ρ=0.3)	**142.36**

Table 6.3 – Summary of Results

While the economic benefits of low-salinity waterflooding are very substantial, it is important to incorporate the value of downstream managerial flexibility to understand more clearly the sources of value. As an example, we can observe in the above summary that the value of the project, subject to the relevant risks but managed optimally, actually exceeds the deterministic success case economics that ignores all risks. Although low-salinity water flooding directly impacts only oil recovery, we see that it indirectly affects the optimal timing of gas production. This is important, as the ability to optimize timing under price uncertainty for both commodities is over 80% of the incremental value to the deterministic case, or $8.6 Bn. The sources of incremental value relative to the base deterministic case without new technology are shown in Figure 6.15.

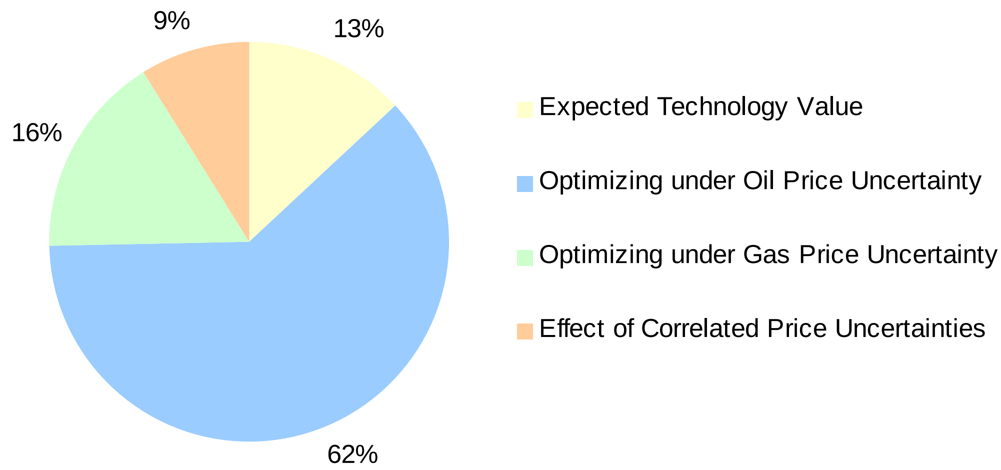

Figure 6.15 – Sources of Incremental Value above Base Case

The use of a dynamic economic model also has implications for the inputs and their estimation. In Figure 6.16 we show a tornado chart of three key inputs we examined for the base case in Figure 6.5. While the ranges in expected value due to the oil and gas price variables were approximately \$75 Bn and \$18 Bn, respectively before, the figure shows these ranges have been reduced to approximately \$15 Bn and \$3 Bn, respectively now.

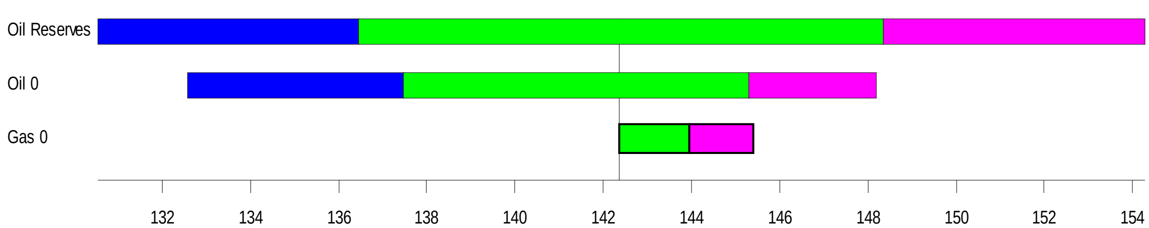

Figure 6.16 – Sensitivity Analysis for Dynamic Model

In contrast, oil reserves uncertainty, which we have not modeled as a stochastic process in this model, has about the same impact on the solution as before. The implication is that the specification of a stochastic process for an uncertainty and its subsequent use in a dynamic economic framework eliminates the requirement for a single point estimate of a variable to impound all of the information about that particular uncertainty. This can be an important change, especially for variables such as commodity price that have high levels of uncertainty in forecasting.

6.5 CONVERGENCE PROPERTIES OF NUMERICAL SOLUTION

The binomial tree model used in the last section is useful for determining bounds and intuitive analysis; however the approximation errors for a one-year time increment are likely to be significant, as was discussed earlier. Even though steps were taken in the previous section to reduce the number of nodes necessary, a reduction of time increments to a length of one half-year would double the number of chance nodes, which exceeds the limit of the size of problem that can be practically solved in a decision tree. Thus, for a large practical problem like the one being considered here, showing convergence and obtaining a more accurate solution is not possible in decision trees. Fortunately, the endpoints in this model formulation are recombining, as was discussed in Section 3. Therefore the convergence properties can be investigated by switching to lattice format.

To switch to lattice format, a different kind of algorithm is required. In this case the algorithmic approach presented in Section 4.2 was implemented in Visual Basic with an Excel interface. With this approach, it is possible to model 200 or more increments within reasonable computational times. The results, shown in Figure 6.17, indicate agreement with the decision tree model for the case of 20 annual periods, and also show downward bias in this first estimate.

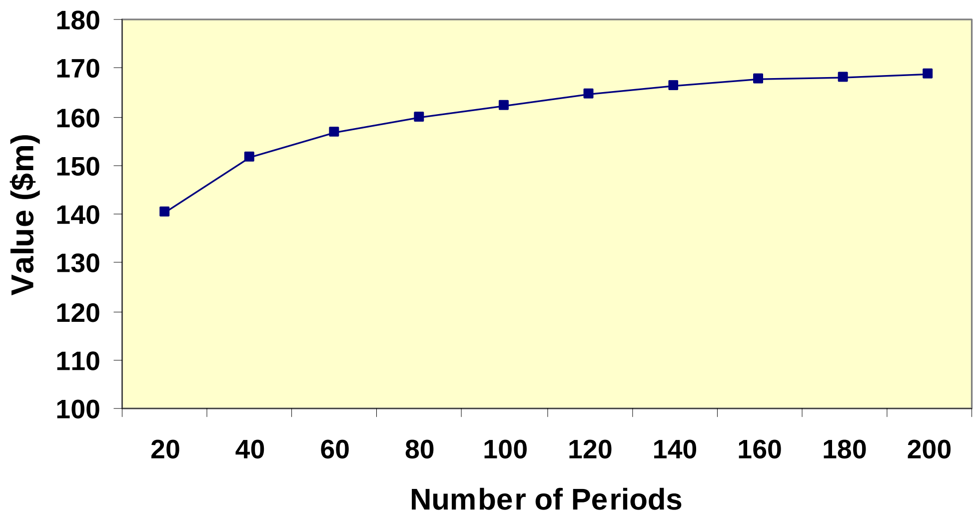

Figure 6.17 – Convergence of Switching Option Value

The convergence behavior shown in the above figure mirrors that obtained by Clewlow and Strickland (2000) for their bivariate binomial approximation. The figure shows that using the binomial approximation with a low level of time granularity can result in significant error, as expected. While the initial estimate with annual time increments in this case was approximately 16% lower than the true solution, which was estimated by extrapolation of the above curve, this bias was reduced to just over 3% by reducing the increment to five periods per year. This would probably be considered a reasonable estimate for most real options applications.

7. CONCLUSIONS

In this dissertation, we have shown how to construct recombining binomial lattices or binomial trees to model underlying stochastic processes that are mean-reverting, and have extended this approach to develop a method for modeling two-factor processes and combined correlated one-factor processes. This method provides an improved approach relative to simulation-based algorithms and discrete trinomial trees, and facilitates the evaluation of real options with early-exercise characteristics, as well as multiple concurrent options.

We have shown how convergence is achieved for this method by reducing the discrete time increments and have described the behavior of models for several example problems. The models developed in this research have been tested by implementing the lattice in binomial decision tree format for small problems, and we have also developed algorithms to implement in lattice format for problems where the number of periods becomes large and beyond the capabilities of commercial decision tree software.

Three different data analysis techniques, Kalman filtering, seemingly unrelated regression, and an implied approach with futures data have been tested for their ability to estimate mean-reverting stochastic process parameters and for their computational requirements. The Kalman filter is a computationally intense approach, but it provides stable estimates of the parameters as well as error estimates, and we were able to replicate parameter estimation work done by other researchers. Although a rigorous statistical comparison of the approaches was not undertaken, we found the implied approach could be used to approximate parameter estimates from the Kalman filter approach on a limited basis in cases where computational burden is a consideration.

In the concluding section, we illustrated a real application by solving for the value of an oil and gas switching option related to a new enhanced oil recovery technology that would be applied to the North Slope Alaska producing area. The value of the technology, the broader value of North Slope oil and gas production, and operating decisions about the Alaska Gas pipeline are all interrelated in this problem. We first considered a base case deterministic model for continued development of the North Slope without the new technology, and then added uncertainties incrementally, starting with the private uncertainties related to the technology. We then added oil and gas price uncertainties by using a binomial approximation of two correlated one-factor mean-reverting models, to finally develop a more fully dynamic economic model of the problem. As would be expected, the solutions from our model were somewhat different from the base case deterministic model, showing the value in making optimal decisions under uncertainty. Although this project has robust economics in all cases, even the base case, it is important to understand and capture all of the underlying sources of value, as the project may be in competition with other high value projects in a constrained capital budgeting environment. The analysis provided here also provides guidance and insight on operating decisions that would not be obtained through a deterministic model. Results from our study could, for example, also be used to inform decisions about construction of the Alaska Gas Pipeline.

Further research issues regarding the methods we have developed here include additional work in the areas of parameter estimation and empirical testing. Stochastic process parameter estimation is obviously critical to formulating a good model, however empirical work in this area is limited. While we investigated this topic to support the models developed for this research, the estimation techniques we used could be the subject of a more detailed statistical comparison.

We have made several assumptions in this work that could be tested in an expanded study. For example, we have assumed that when a switch in production mode is made, it occurs instantaneously. In reality, it would take an estimated two years to convert the wells and bring the pipeline into operation. This lag time between decision and operational change could be factored into the model by continuing to model uncertainty past the point of decision for two additional periods for each decision node, and possibly including a decision to delay actual execution of the switch if conditions worsen during this period.

The approach we have developed here could also be tested for financial options if a suitable application can be found. In such a case, solutions from alternative methods might be available and could be compared with solutions from our method. The primary alternative to our method is the trinomial tree approach of Hull and White (1994b), however here has been very little published work done to empirically test their approach for two-factor processes. Our approach could be tested in parallel with the Hull-White model to compare the accuracy and computational requirements of the two different methods.

REFERENCES

Amin, K., "On the Computation of Continuous Time Option Prices Using Discrete Approximations," *The Journal of Financial and Quantitative Analysis,* **26** (1991), 477-495.

Amin, K., "Jump-Diffusion Option Valuation in Discrete Time," *The Journal of Financial,* **48** (1993), 1833-1863.

Amin, K. and Khanna, A., "Convergence of American Option Values from Discrete to Continuous Time Financial Models," *Mathematical Finance,* **4** (1994), 289-304.

Amin, K. and Bodurtha, J., "Discrete-Time Valuation of American Options with Stochastic Interest Rates," *The Review of Financial Studies,* **8** (1995), 193-234.

Bailey, W., Bhandari, A., Faiz, S., Srinivasan, S., and Weeds, H., "Unlocking the Value of Real Options," *Oilfield Review*, **15** (2003) 4-19.

Bessembinder, H., Coughenour, J., Sequin, P., and Smoller, M., "Mean Reversion in Equilibrium Asset Prices: Evidence from Futures Term Structure," *Journal of Finance,* **50** (1995), 361-375.

Bhattacharya, S., “Project Valuation with Mean-Reverting Cash Flow Streams,” *The Journal of Finance*, **33** (1978), 1317-1331.

Black, F., “The Pricing of Commodity Contracts,” *Journal of Financial Economics*, **3** (1976), 167-179.

Black, F. and Scholes, M., “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, **81** (1973), 637-59.

Boyle, P., “A Lattice Framework for Option Pricing with Two State Variables,” *Journal of Financial and Quantitative Analysis*, **23** (1988), 1-12.

Boyle, P., Evnine, J., and Gibbs, S., “Numerical Evaluation of Multivariate Contingent Claims,” *The Review of Financial Studies*, **2** (1989), 241-250.

Boyle, P. and Tse, Y., “An Algorithm for Computing Values of Options on the Maximum or Minimum of Several Assets,” *The Journal of Financial and Quantitative Analysis*, **25** (1990), 215-227.

Brandao, L. and Dyer, J., “Decision Analysis and Real Options: A Discrete Time Approach to Real Option Valuation,” Forthcoming in the *Special Volume of Annals of Operations Research – Recent Advances in Decision Making under Uncertainty*, 2004.

Brandao, L., Dyer, J., and Hahn, W., "Using Decision Analysis to Solve Real Option Valuation Problems – Part 1: Building a Generalized Approach," Working Paper, The University of Texas at Austin (2004).

Brennan, M. and Schwartz, E., "Finite Difference Methods and Jump Processes Arising in the Pricing of Contingent Claims: A Synthesis," *The Journal of Financial and Quantitative Analysis,* **13** (1978), 461-474.

Brennan, M., "The Pricing of Contingent Claims in Discrete Time Models," *The Journal of Finance*, **34** (1979), 53-68.

Brennan, M. and Schwartz, E., "Evaluating Natural Resource Investments," *The Journal of Business,* **58** (1985), 135-157.

Calistrate, D., Paulhus, M., and Sick, G., "A Recombining Binomial Tree for Valuing Real Options with Complex Structures," Working Paper, The University of Calgary (1999).

Carlson, M., Khokher, Z., and Titman, S., "An Equilibrium Analysis of Exhaustible Resource Investments," Working Paper, The University of Texas at Austin (2002).

Carr, P., "The Valuation of Sequential Exchange Opportunities," *Journal of Finance,* **43** (1988), 1235-1256.

Childs, P., Ott, S., and Triantis, A., "Capital Budgeting for Interrelated Projects: A Real Options Approach," *Journal of Financial and Quantitative Analysis*, **33** (1998), 305-334.

Clewlow, L. and Stickland, C., *Implementing Derivatives Models*, Wiley, West Sussex, U.K., 2000.

Copeland, T. and Antikarov, V., *Real Options*, Texere LLC, New York, 2001.

Cox, J. and Ross, S., "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, **3** (1976), 145-166.

Cox, J., Ross, S. and Rubinstein, M., "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, **7** (1979), 229-263.

Cox, J., Ingersoll, J., and Ross, S., "The Relationship Between Forward Prices and Futures Prices," *Journal of Financial Economics*, **9** (1981), 321-346.

Cox, J., Ingersoll, J., and Ross, S., "A Theory of the Term Structure of Interest Rates," *Econometrica*, **53** (1985), 385-467.

Dias, M., Rocha, K., and Teixeira, J., "The Optimal Investment Scale and Timing: A Real Option Approach to Oilfield Development," Presented at the 8th Annual Real

Options Conference, Montreal, Canada (2003). Available for download at http://www.realoptions.org/

Dixit, A., "Entry and Exit Decisions Under Uncertainty," *The Journal of Political Economy*, **97** (1989), 620-638.

Dixit, A. and Pindyk, R., *Investment Under Uncertainty*, Princeton University Press, Princeton, NJ, 1994.

Geske, R. and Johnson, H., "The American Put Option Valued Analytically," *Journal of Finance*, **39** (1984), 1511-1524.

Geske, R. and Shastri, K., "Valuation by Approximation: A Comparison of Option Valuation Techniques," *Journal of Financial and Quantitative Analysis*, **20** (1985), 45-71.

Gibson, R. and Schwartz, E., "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims," *Journal of Finance*, **45** (1990), 959-976.

Gundy, B., "Option Prices and the Underlying Asset's Return Distribution," *Journal of Finance*, **46** (1991), 1046-1069.

Harvey, A., *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press, Cambridge, U.K., 1989.

He, H., "Convergence from Discrete to Continuous Time Contingent Claims Prices," *Review of Financial Studies*, **3** (1990), 523-546.

Hilliard, J. and Reis, J., "Valuation of Commodity and Futures and Options Under Stochastic Convenience Yields, Interest Rates, and Jump Diffusions in the Spot," *Journal of Financial and Quantitative Analysis*, **33** (1998), 61-86.

Hjalmarsson, E., "Does the Black-Scholes Formula Work for Electricity Markets? A Non-Parametric Approach," Working Paper, Yale University (2003).

Ho, T., Stapleton, R., and Subrahmanyam, M., "Multivariate Binomial Approximations for Asset Prices with Non-Stationary Variance and Covariance Characteristics," *Review of Financial Studies*, **8** (1995), 1125-1152.

Howard, R., "Options," in Zeckhauser, R., Keeny, R., and Sebenius, J., Eds., *Wise Choices*, Harvard Business School Press, Boston, MA, 81-101 (1996).

Hull, J. and White, A., "Pricing Interest Rate Derivative Securities," *Review of Financial Studies*, **3** (1990a), 573-592.

Hull, J. and White, A., “Valuing Derivative Securities Using the Explicit Finite Difference Method,” *The Journal of Financial and Quantitative Analysis,* **25** (1990b), 87-100.

Hull, J. and White, A., “One-Factor Interest Rate Models and the Valuation of Interest Rate Derivative Securities,” *The Journal of Financial and Quantitative Analysis,* **28** (1993a), 234-254.

Hull, J. and White, A., “Efficient Procedures for Valuing European and American Path-Dependent Derivatives,” *Journal of Derivatives,* **1** (1993b), 21-31.

Hull, J. and White, A., “Numerical Procedures for Implementing Term Structure Models I: Single-Factor Models,” *Journal of Derivatives,* **2** (1994a), 7-16.

Hull, J. and White, A., “Numerical Procedures for Implementing Term Structure Models II: Two-Factor Models,” *Journal of Derivatives,* **2** (1994b), 37-48.

Hull, J. and White, A., “Using Hull-White Interest Rate Trees,” *Journal of Derivatives,* **4** (1996), 26-36.

Hull, J., *Options, Futures and Other Derivatives,* Prentice Hall, New Jersey, 1999.

Jaillet, P., Ronn, E., and Tompaidis, S., “ Valuation of Commodity-Based Swing Options,” *Management Science,* **50** (2004), 909-921.

Kalman, R., “A New Approach to Linear Filtering and Prediction Problems,” *Transactions of the ASME,* **82** (1960), 35-45.

Laughton, D. and Jacoby, H., “Reversion, Timing Options, and Long-Term Decision Making,” *Financial Management,* **33** (1993), 225-240.

Lari-Lavassani, A., Simchi, M, and Ware, A., “A Discrete Valuation of Swing Options”, *Canadian Applied Mathematics Quarterly*, **9** (2001) 35-73.

Lo, A. and Wang, J., “Implementing Option Pricing Models When Asset Returns Are Predictable,” *Journal of Finance,* **50** (1995), 87-129.

Longstaff, F. and Schwartz, E., “Valuing American Options by Simulation: A Simple Least-Squares Approach,” *The Review of Financial Studies*, **14** (2001), 113-147.

Madan, D., Milne, F., and Shefrin, H., “The Multinomial Option Pricing Model and Its Brownian and Poisson Limits,” *The Review of Financial Studies,* **2** (1989), 251-265.

Margrabe, W., “The Value of an Option to Exchange One Asset for Another,” *The Journal of Finance*, **33** (1978), 177-186.

McGuire, P., Chatham, J., Paskvan, F., Sommer, D., and Carini, F., "Low Salinity Oil Recovery: An Exciting New EOR Opportunity for Alaska's North Slope," SPE Paper #93903, Presented at the SPE Western Regional Meeting, Irvine, CA (2005).

Merton, R., "The Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, **4** (1973), 141-183.

Metcalf, G. and Hassett, K., "Investment Under Alternative Return Assumptions: Comparing Random Walks and Mean Reversion," Technical Working Paper No. 175, National Bureau of Economic Research, Cambridge, MA (1995).

Muck, M. and Rudolf, M., "An Extension of the Discrete Implementation of the Hull and White Two-Factor Model," Working Paper, WHU (2002).

Nau, R. and McCardle, K., "Arbitrage, Rationality, and Equilibrium," *Theory and Decision*, **31** (1991), 199-240.

Nelson, D., "ARCH Models as Diffusion Approximations," *Journal of Econometrics*, **45** (1990), 7-38.

Nelson, D. and Ramaswamy, K., "Simple Binomial Processes as Diffusion Models in Financial Models," *The Review of Financial Studies*, **3** (1990), 393-430.

Peterson, S., Stapleton, R., and Subrahmanyam, M., "The Valuation of American-Style Swaptions in a Two-Factor Spot-Futures Model," Working Paper, New York University (1999).

Pindyck, R. and Rotemberg, J., "The Excess Co-Movement of Commodity Prices," *The Economic Journal*, **100** (1990), 1173-1189.

Ramaswamy, K. and Sundaresan, S., "The Valuation of Options on Futures Contracts," *The Journal of Finance*, **40** (1985), 1319-1340.

Ribeiro, D. and Hodges, S., "A Two-Factor Model for Commodity Prices and Futures Valuation," Presented at the European Financial Management Association Annual Meeting, Basel, Switzerland (2004).

Robel, G., "Real Options and Mean-Reverting Prices," Presented at the 5th Annual Real Options Conference, Los Angeles, CA (2001). Available for download at http://www.realoptions.org/

Ronn, E., "The Valuation of Oil Fields as Optimal Exercise of the Extraction Option," Working Paper, The University of Texas at Austin (2004).

Rubinstein, M., "Implied Binomial Trees," *The Journal of Finance*, **49** (1994), 771-818.

Rubinstein, M., “The Valuation of Uncertain Income Streams and the Pricing of Options,” *Bell Journal of Economics,* **7** (1976), 407-425.

Schwartz, E., “The Valuation of Warrants: Implementing a New Approach,” *Journal of Financial Economics*, **4** (1977), 79-94.

Schwartz, E., “The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging,” *Journal of Finance*, **52** (1997), 923-973.

Schwartz, E., “Valuing Long-Term Commodity Assets,” *Financial Management*, **27** (1998), 57-66.

Schwartz, E. and Cortazar, G., “The Valuation of Commodity-Contingent Claims,” *The Journal of Derivatives*, **1** (1994), 27-39.

Schwartz, E. and Smith, J., “Short-Term Variations and Long-Term Dynamics in Commodity Prices,” *Management Science*, **46** (2000), 893-911.

Sharpe, W., *Investments*, Prentice Hall, New Jersey, 1978.

Slade, M., “Valuing Managerial Flexibility: An Application of Real Option Theory to Mining Investments,” *Journal of Environmental Economics and Management*, **41** (2001), 193-233.

Smith, J. and McCardle, K., “Valuing Oil Properties: Integrating Option Pricing and Decision Analysis Approaches,” *Operations Research*, **46** (1998), 198-217.

Smith, J. and McCardle K., “Options in the Real World: Lessons Learned in Evaluating Oil and Gas Investments,” *Operations Research*, **47** (1999), 1-15.

Smith, J. and Nau, R., “Valuing Risky Projects: Option Pricing Theory and Decision Analysis,” *Management Science*, **41** (1995), 795-816.

Staley, M. and Wicentowich, C., “Spline Trees: Binomial Lattices for Multi-Factor Processes,” Working Paper, The University of Ontario (2003).

Stapleton, R. and Subrahmanyam, M., “The Valuation of Options When Asset Returns are Generated by a Binomial Process,” *The Journal of Finance*, **5** (1984), 1525-1539.

Stroock, D. and Varadhan, S., *Multidimensional Diffusion Processes*, Springer, Berlin, 1979.

Tourinho, O., "The Valuation of Reserves of Natural Resources: An Option Pricing Approach," Ph.D. Dissertation, University of California, Berkeley (1979).

Trigeorgis, L., *Real Options – Managerial Flexibility and Strategy in Resource Allocation*, MIT Press, 1996.

Vasicek, O., "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, **5** (1977), 177-188.

Webb, K., Black, C., and Al-Ajeel, H., "Low Salinity Oil Recovery – Log-Inject-Log," SPE Paper #89379, Presented at the SPE/DOE Fourteenth Symposium on Enhanced Oil Recovery, Tulas, OK (2004).

VITA

Warren Joseph Hahn was born in Fredericksburg, Texas, on April 12, 1964, the son of Warren Julius Hahn and Martha Kate Hahn. After graduating from Fredericksburg High School in 1982, he entered The University of Texas at Austin. He received Bachelor of Science degrees in Petroleum Engineering and Civil Engineering in May 1988 and December 1989, respectively, and a Master of Science in Environmental Engineering in May 1991. From July 1991 trough August 2001, he worked for BP America and predecessor companies in various engineering and financial management roles. He re-entered The University of Texas in August 1999 and received a Masters of Business Administration in May 2001. In September of 2001, he entered the Graduate School of Business at The University of Texas as a Ph.D. student in Management Science and Information Systems.

Permanent Address: 7613 Ranch Road 648

Doss, Texas 78618

This dissertation was typed by the author.