The Pennsylvania State University

The Graduate School

Department of Economics

AN EMPIRICAL INVESTIGATION OF INVESTMENT UNDER UNCERTAINTY WITH SUNK COSTS: OIL PRODUCTION IN OKLAHOMA

A Thesis in

Economics

by

Boris Molls

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2001

We approve the thesis of Boris Molls.

Date of Signature

Mark J. Roberts
Professor of Economics
Thesis Advisor
Chair of Committee

Timothy J. Considine
Professor of Energy, Environmental, and
Mineral Economics

Susumu Imai
Assistant Professor of Economics

James R. Tybout
Professor of Economics

James Jordan
Professor of Economics
Acting Head of the Department of Economics

ABSTRACT

We observe a very distinct pattern of entry and exit in the oil industry. Significant numbers of well openings and closings are only observed during periods of very high oil prices and very low oil prices, respectively. A band of inaction lies between the high entry and the low exit price. Recent economic theory finds that the combination of sunk costs and uncertainty of future market conditions leads to this pattern of hysteresis.

This dissertation develops an empirical model of firm investment that can measure the role of sunk costs and uncertainty on investment and production. In my dynamic discrete choice model firms make the decision to produce oil, mothball, or shutdown oil wells. Switching between the states production, mothballing, and shutdown involves non-recoverable cost. The firm's decision is affected by current and expected future market conditions, including the price of oil, oil price volatility, and the well's characteristics. The model, which allows me to test if sunk costs and price uncertainty are determinants of the investment choice, is used to explain the supply response of oil producers to fluctuations in the market price of oil.

The model is estimated using a recently developed panel-probit simulation estimator that allows me to control for several sources of correlation in the long well history data, including persistent time-invariant random effects and serial correlation. Using a new micro-level data set on oil production units in Oklahoma it is found that sunk costs are an important determinant of the production status. Although the probability of being a producer increases in the price of oil, oil producers only slowly respond to price changes because of non-recoverable switching costs. Moreover, oil price volatility affects the decision to change the production status.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This dissertation was influenced and improved by advice, knowledge, and assistance of several people. First and foremost, my thanks go to my dissertation advisor Mark Roberts for his guidance, support, and encouragement. His help was invaluable for completing and improving this project. I am grateful to my committee members Timothy Considine, Susumu Imai, and James Tybout for their help and very useful comments.

This research was only possible with the data provided by Sarkeys Energy Center at the University of Oklahoma. Kathryn E. Hines of the Sarkeys Energy Center provided priceless assistance answering my questions regarding the data. My appreciation extends to Timothy Dunne for helping to acquire the data set.

This research was supported by NSF grant SBR-9631766.

To Jasmina, Michael, and Roshni.

Chapter 1

INTRODUCTION

Within just the last two years the price of a barrel of crude oil has varied from a low of $11 to more than $35 and this has had a substantial impact on the profitability of individual U.S. oil wells. At the high prices, even the least productive wells are quite profitable, yet oil producers are reluctant to drill new wells. During the periods of low prices, many wells cannot cover their variable costs of production, yet oil companies are slow to abandon these unprofitable wells. In general, large numbers of well openings are only observed when the price of oil is very high and large numbers of well closing are only observed when the price of oil is very low. Between the two prices that trigger entry and exit is a band that generates little well drilling or plugging activity. This pattern, in which oil well opening and closing does not respond quickly or smoothly to changes in the world price of oil, is important for understanding the domestic crude oil supply.

Recent economic theory has found that the combination of sunk costs and uncertainty in future market conditions is capable of generating the sluggish or lumpy patterns of investment observed in many industries. For instance, Dixit (1989) analyses a firm's entry and exit decision when future output prices are uncertain and market entry and exit require sunk costs. In his model, the optimal investment strategy is a pair of entry and exit trigger prices, where the exit threshold is lower than the entry threshold. Between the two thresholds is a band of inaction, referred to as the hysteresis band, in which price movements do not lead to any investment activity by the firm.[1] The entry trigger price is increasing, and the exit trigger price is decreasing, in sunk costs and uncertainty. Intuitively, uncertainty gives firms an incentive to delay investment and wait for new information to arrive before sinking entry and exit costs. This model can be extended in various ways. For example, to a model of entry and exit with temporary shutdown, called mothballing, as in Dixit and Pindyck (1994).

In the oil industry there is substantial uncertainty of future market prices and significant sunk costs of well investment. Hence, there are incentives for firms to delay seemingly profitable investments, such as opening new wells even when oil prices rise. The same forces give firms an incentive to continue unprofitable activities, such as keeping wells open in periods of low prices, in the expectation that future market conditions will improve. In order to understand the supply response of the domestic U.S oil industry to fluctuations in market price, this dissertation develops an empirical, discrete choice model of a firm's decision to produce oil, mothball, or shutdown oil wells. Current and expected future profits of the firm depend on the current production status so that sunk costs and uncertainty affect the firm's investment choices. Using a dynamic program, the model illustrates under which conditions the empirically observable choices are made and derives an empirical estimation equation.

The model, which allows testing if sunk costs and price uncertainty are determinants of the investment choice, is estimated using a micro-level data set for oil wells and oil leases in Oklahoma. This new and unique data set allows me to see which firms exercise their option to switch into another production state, and which do not. In addition, the data allows sunk cost to differ across individual observations. The empirical model is used to explain the supply response of oil producers to fluctuations in the market price of oil.

[1] "Hysteresis" is the tendency for an effect to persist after the cause that brought it about has disappeared.

Within the theoretical framework of my oil production model it is shown that the lagged production status matters for the current choice if there are sunk costs. Thus, testing for state dependence, it is possible to test if sunk costs have an impact on the decision to produce, mothball, or shutdown. Two models are estimated: a two-choice model that distinguishes between the state production and no-production, and a three-choice model that allows firms to produce, mothball, or shutdown. The shutdown and mothballing states are a further refinement of the no-production state. The discrete choice model is estimated with a dynamic multinomial probit estimator. Using a simulation estimator developed by Keane (1994), I estimate the model allowing for a firm level profit shock, and serial correlation in profit shocks.

The estimation results show that there is persistence in the states no-production, mothballing, and shutdown. The current production status is significantly affected by the production status in the previous period. Controlling for other possible sources of profit persistence, this phenomenon is attributed to the importance of sunk costs. Moreover, it is found that state persistence is increasing in a well's depth, which serves as a proxy for sunk costs.

This dissertation is structure in the following way: Chapter 2 reviews the relevant theoretical sunk costs and real options literature and discusses earlier empirical findings. Chapter 3 gives an overview of important engineering and economic factors involved in oil production. Moreover, it presents aggregate data from Oklahoma and Texas that provides first evidence that sunk costs play an important role in the oil industry. Chapter 4 presents a model oil production and derives its reduced form estimation equation. Chapter 5 describes the construction of the new well and lease data set from Oklahoma that is used for estimating the model and presents statistics of the data. Finally, in Chapter 6 the oil production model is estimated and the estimation results are discussed. Chapter 7 concludes the dissertation.

Chapter 2

REVIEW OF THE LITERATURE

This chapter reviews the relevant theoretical and empirical literature. Section 2.2 highlights and summarizes important theoretical findings. Section 2.3 discusses how empirical research attempts to test sunk costs theories and presents the empirical literature's findings. Finally, this dissertation is placed in the existing literature.

2.1 Introduction

Most investment decisions have three underlying features. First, investment is irreversible. Some costs of investment cannot be recovered and are sunk. Second, the environment in which the investment is undertaken changes over time. Investment is risky and new information arrives as time passes. Third, most investment opportunities do not vanish at once. The investment decision does not only invoke the question whether investment should be undertaken or not, it also involves the aspect of when to invest.

Sunk costs, uncertainty, and persistence together give value to an investment opportunity itself. One can view an investment opportunity as an option to invest. The investment option can be exercised today, or sometime in the future. In fact, real investment options, for instance the opportunity to build a new plant, are analogous to American call options on common stock. Exercising the real option involves sinking investment costs and exercising the stock option implies spending the strike price. In each case, upon exercising, the investor receives an asset that has a current value along with the claim to an uncertain future profit stream. Nevertheless, the option does not need to be exercised and, thus, bad economic states of the world can be avoided. This gives value to the option itself.

Most investment models do not take the options value in to account when analyzing investment. For instance, consider the net present value rule: here a firm should invest if the net present value (NPV) of the investment project is larger than the investment costs. This investment rule does not account for the fact that investment implies losing the investment option. After accounting for the option value, with its opportunity costs of exercising the option, the investment rule should be as follows: invest if the NPV is larger than investment costs plus the value of the investment option. The same idea can easily be extended to a firm's exit decision: a firm should scrap its investment if the NPV of keeping the investment plus the scrapping option are worth less than the scrapping costs.

Theoretical sunk cost models of investment find that the option value can be substantial. Consequently, the NPV threshold trigger for entry is considerably higher when taking the option value into account. Similarly, when analyzing exit behavior, the exit threshold is lower when taking the value of the exit option into account. As a result, sunk cost models of investment are able to explain entry and exit hysteresis, and in general the sluggish and lumpy patterns of investment observed in many industries. Older investment models, in contrast, disregard the option value by assuming investment is largely reversible, investment opportunities are once-and-for-all, or investment can be undertaken in marginal increments.[2] Accordingly, an important empirical research question is whether sunk costs are important for explaining investment

[2] Contradicting the assumption of investment in marginal increments, there is strong evidence that speaks in favor of lumpy investment, see Doms and Dunne (1998).

behavior? If so, how important are they? This question is especially of interest in the light of the fact that older models of investment have largely failed to explain and predict investment.[3]

Subsection 2.2 reviews some of the theoretical literature and subsection 2.3 reviews the relevant empirical literature.

2.2 Theory

Recent economic theory finds that the combination of sunk costs and uncertainty of future market conditions is capable of generating entry and exit hysteresis. These sunk cost models of investment can explain the sluggish or lumpy patterns of investment observed in many industries. The foundation for the findings in this literature is based in the fact that the combination of sunk costs and uncertainty gives value to an investment opportunity itself, driving up the opportunity costs of investment.

Early studies of irreversible investment, such as Arrow and Fisher (1974), Henry (1974), or Bernanke (1983), find that it is important to consider irreversibility and uncertainty when analyzing investment decisions. If investment is irreversible, even the optimal investment decision of a risk-neutral investor is affected by uncertainty. Nevertheless, these early studies do not derive an explicit investment rule, or an explicit formula for option value and the value of investment.

---

[3] Dixit and Pindyck (1994, p. 419) conclude: “The explanation of aggregate and sectoral investment spending has been one of the less successful endeavors in empirical economics. For the most part, econometric models have not been very useful for explaining or predicting investment spending. The problem is not just that these models have been unable to explain and predict more than a small portion of the movements of investment. In addition, constructed quantities that in theory should have strong explanatory power - such as Tobin’s q, or various measures of the cost of capital – in practice do not.”

McDonald and Siegel (1986) introduce financial option pricing techniques to the analysis of real investment, making it feasible to derive an explicit investment rule and an explicit formula for the value of the option to invest. Their paper analyzes the decision of a risk-neutral monopoly to obtain an investment project with uncertain future value. Exercising the investment option involves sunk costs and future values of the investment project are governed by a geometric Brownian motion. McDonald and Siegel show that a firm invests if the value of its investment project passes a threshold. Moreover, the investment threshold and the value of the investment option are increasing in sunk costs and uncertainty over future project values. Intuitively, since investment costs are sunk and the investment value process follows a geometric Brownian motion a firm only invests if the value of the investment project is high and unlikely to be very low any time soon. In this model it can be optimal for the firm to postpone investment until benefits from investment are twice the investment costs. Although this paper considers only a quite restricted case, it sets the framework for later studies of irreversible investment under uncertainty.

Dixit and Pindyck (1994, Chapter 5) relate the results from the McDonald-Siegel model to two neo-classical investment rules. [4] The Marshallian investment rule says to invest when the value of the expected profits is larger than investment costs. The Jorgenson investment threshold is at the point where the marginal profit from an extra unit of capital is equal to the user costs of capital. Dixit and Pindyck show that in either model the investment threshold is lower than what it should be according to the McDonald-Siegel model. The difference arises because the neo-classical investment models do not take the option value of investment into account. The higher

---

[4] A very comprehensive overview and summary of the literature is Dixit and Pindyck (1994), *"Investment Under Uncertainty."* For the most part the models in this book build on the McDonald-Siegel model. Dixit and Pindyck explain the basic model and its mathematics, and review extensions that can be found in the literature. The extensions cover topics such as: dynamic industry equilibrium in a competitive industry, sequential investments, investment and capacity choice, and policy interventions when investment is irreversible. The basic results

the underlying uncertainty or sunk costs, the larger is the difference between the proposed investment rules.

Dixit (1989) analyses the market entry and exit decision of a risk neutral firm when entry and exit involve sunk costs. If a firm is idle it can invest into a single discrete project that produces one unit of output per period. An active firm can scrap the project. Uncertainty enters the model in the form of uncertainty over future output prices. The output price, and hence the investment project value, is governed by a geometric Brownian motion. The idle firm has the option to become an active firm. Exercising this option can be optimal if the output price is larger than the variable costs of production. The active firm has the option to become an idle firm and exercising this option can be optimal if the output price is below the variable costs of production. Dixit finds that the optimal investment strategy is a pair of entry and exit trigger prices. The idle firm invests if the output price surpasses an entry threshold and the active firm closes down if the output price falls below an exit threshold. The exit threshold is lower than the entry threshold, which implies that there is there is a band of investment inactivity between the two thresholds. Moreover, the entry trigger price is increasing in sunk costs and uncertainty, and the exit trigger price is decreasing in sunk costs and uncertainty. The intuition behind these results is the same as in the other papers discussed above. The option value of switching into another state is increasing in sunk costs and uncertainty. The higher sunk costs or uncertainty are, the higher the opportunity costs of switching into the alternative state. Thus, the idle firm has an incentive to invest only if investment looks very profitable, resulting in a high entry threshold, and the active firm has an incentive to scrap its investment only if the future looks very bad, resulting in a low exit threshold. Uncertainty gives a firm an incentive not to change its production status and, instead, wait for new information before committing to a new production

---

regarding, for example entry and exit hysteresis along with its comparative static results regarding state. One of the important implications of these findings is that a firm's current production status does not only depend on the current market conditions, but also on its production status in the previous time period as well. A firm's investment history matters for its current production status.

Dixit (1988) is one of the many possible extensions to the McDonald and Siegel model of irreversible investment. A similar extension is Dixit and Pindyck (1994, Chapter 7), which is an investigation of entry and exit with temporary shutdown also called mothballing. The model environment is the same as in Dixit (1988) and the investment options are similar. A firm without production capital in place has the option to invest and install production capacity. A firm with capital in place has the option to scrap its capital, to produce, or to mothball its production capacity. The last option is an addition to the Dixit (1988) model. Mothballing, an intermediate state between production and complete shutdown, gives the firm the opportunity to stop production without losing its capital. However, there are fixed costs of switching between mothballing and production, and mothballing involves per period mothballing costs. In this world mothballing can be a viable strategy if the output price is lower than the variable costs of production. If the output price is expected to remain below the variable costs of production for a long enough time it may be optimal to scrap the investment project altogether.

Solving the model, Dixit and Pindyck find that there is an entry (initial investment) and an exit (scrapping) threshold. In between the entry and the exit thresholds is a mothballing region with a production re-entry threshold and a mothballing threshold. A producing firm mothballs at some low price and a mothballed firm may re-enter when the price is higher again. In essence, there is a mothballing hysteresis band embedded into an entry and exit hysteresis band. The comparative statics of this model are similar to the ones in the previous model. The upper

---

sunk costs and uncertainty, are very robust.

threshold of each hysteresis band is increasing and the lower threshold is decreasing in uncertainty and sunk costs. Thus, investment and production history matter for the current status of the firm.

Sunk costs and irreversibility are one way of introducing frictions into investment models[5]. Another way of introducing frictions is to assume that there are adjustment costs of changing the capital stock. Abel and Eberly (1994) try to close the gap between the adjustment cost and irreversible investment literature, and analyze investment under uncertainty when there exist both types of frictions: irreversibility and adjustment costs. They find that investment is a non-decreasing function of the shadow value of installed capital. Furthermore, with both types of frictions gross investment falls into three regimes: a region of positive investment, a region of inactivity, and a region of negative investment. However, if adjustment costs are the only friction and the investment cost function is differentiable at the point where there is no investment, the region of investment inactivity disappears. In addition, if the dis-investment adjustment costs are large enough, it is never optimal for the firm to undertake negative gross investment and investment may observationally appear to be completely irreversible.

To sum up, theoretical models of investment that take the value of investment opportunities into account are able to generate highly non-linear investment patterns. Along with current and future expected market conditions, a firm's investment and production history matters for explaining a firm current production status when the costs of switching from one production state to another state are sunk. The production behavior of producing firms is very different from non-producing firm, who need to invest ("enter") first for being able to produce.

[5] Jorgenson (1963) shows that the firm's investment decision is essentially static and equates the marginal product of capital to the user costs of capital if the capital stock can be adjusted instantaneously and costlessly. The investment decision only becomes an interesting dynamic problem if there are frictions.

Theory suggests that sunk costs and uncertainty create entry and exit hysteresis. Sunk costs and uncertainty appear to be important for understanding market dynamics. It may be very important to understand investment and production decisions at the firm level in order to understanding aggregate production behavior. Although Dunn, Roberts, and Samuelson (1988, 1989) document the importance of entry and exit for industry dynamics using firm level data, there have been astonishingly few attempts to empirically investigate the relation between firm investment and sunk costs on the micro-level.

2.3 Empirical Literature

Based on theoretical results from the sunk costs literature one can derive predictions for empirically observable variables. For example, it is possible to derive hypotheses regarding the expected firm entry and exit behavior, firm values, and the numbers of entering and exiting firms in a market. Several empirical studies that test the implications of models of investment under uncertainty with sunk costs are summarized below.

One way of examining the real options model is to use its predictions regarding investment project values. Valuing a potential project by just using its net present value does not account for the option value. Thus, project values based on the NPV rule are too low. Paddock, Siegel, and Smith (1988) base their analysis on this theoretical finding. Building on the McDonald-Siegel model they develop a formula to compute the value of off-shore oil and gas track leases. Then they estimate the values of specific off-shore tracks using real options pricing and discounted cash-flow techniques. Comparing their estimates to the winning bids in oil and gas track auctions they find that the real options values are better predictors of the winning bid. Yet, they do not statistically test one theory against the other. Moreover, their results could be driven by the winners-curse.

Lambson and Jensen (1995; 1998) examine an implication of the real options model for the maximum and minimum observable firm value in an industry. The real options model predicts that the entry and exit triggering thresholds are increasing and decreasing, respectively, in sunk costs. They argue that in equilibrium the highest observable firm value in a particular industry is bounded by potential entry, e.g. the entry threshold, and the lowest observable firm value is bounded by the exit threshold. Accordingly, Lambson and Jensen argue that the difference between the highest and the lowest firm value over time depends on the industry specific sunk costs. Hence, assuming the existence of external profit shocks, a high sunk cost industry will display higher firm value variability over time than a low sunk cost industry.

Lambson and Jenson (1995) use several years of survey data of farm-land values, categorized by crops and regions, from California and Florida. Assuming orchard crop production entails higher sunk costs than annual crop production, they regress measures of land value variability over time on an orchard dummy variable. As expected they find that the value variability is higher for orchard farms. In a later study, Lambson and Jenson (1998) investigate the same question using a data set of publicly trade firms. In this study, inter-temporal firm value variability is based on stock prices and publicly available debt positions. The book value of property, plant, and equipment serves as a sunk cost proxy. Regressing firm value variability on the sunk cost proxy, while controlling for firm size, yields a significantly positive coefficient. Testing for the importance of sunk cost based on their implication for firm value variability, these two papers provide further evidence that sunk costs are an important determinant of firm entry and exit decisions.

Campa (1993) uses the relation between sunk costs, uncertainty, and entry to examine investment. Using market level data, he studies the number of foreign firms entering the U.S. wholesale trade market in the period 1981 to 1987. The real options model applied to this investment problem predicts that the number of entries is positively correlated with the exchange rate, and that it is negatively correlated with exchange rate uncertainty and sunk costs. Campa constructs measures of industry sunk costs and a country-weighted effective exchange rate. Using tobit regressions he finds that a higher exchange rate increases the number of entries, and that larger sunk costs and higher exchange rate volatility decrease the number of entries. The results are interpreted as evidence in support of the real options investment theory.

Hurn and Wright (1994) try to test the implication of the entry-threshold using micro-data. They have data for 108 North Sea oil field developments and study a two-stage oil development process. Based on the real option model's implication that the entry threshold is increasing in uncertainty, they argue that higher oil price volatility increases the incentive to wait for new oil price information before sinking the costs required for the second investment stage. They interpret the real options model as higher uncertainty delays investment. Hurn and Wright test for this using a hazard rate model. They find that a higher oil price "speeds up" investment, but oil price variance is insignificant.

A general problem in empirical sunk costs studies is the unavailability of sunk cost data. For example, the studies discussed above that include sunk costs in their regression rely on sunk costs proxies. Bresnahan and Reiss (1994) develop another method of testing whether sunk costs matter for investment decisions. Since the entry and exit hysteresis band is increasing in sunk costs, one can use the difference between the entry and exit threshold as a measure of the importance of sunk costs. They estimate the unobservable thresholds by relating changes in market structure to changes in market demand, e.g. market size. Since the hysteresis band is increasing in sunk costs, it will take large changes in market demand to change the number of firms in a market if sunk costs are large. Bresnahan and Reiss use data that describes the number of dentists and demand for dental work in distinct geographic areas for two different years. As expected, they find that there are substantial sunk costs.

Roberts and Tybout (1997) test for sunk cost hysteresis in a similar way. They develop an empirical test for hysteresis by asking whether a firm's history matters for the current investment decision. Note that an incumbent has already sunk the investment costs, whereas a potential entrant has not. This makes the market participation decisions of incumbents and potential entrants asymmetric. For example, output prices that are sufficient to keep an incumbent in the market may not be sufficient to induce entry. Thus, hysteresis implies that the observed investment trajectories are affected by a firm's history. Roberts and Tybout show how the researcher can test for the importance of sunk costs by asking whether market participation history matters. This period's market participation depends on last period's market participation only if sunk costs matter. Roberts and Tybout derive this result using a dynamic program that describes a firm's export entry and exit decision. Using a micro-level data set of exporting firms from Columbia they estimate the reduced form of this model. They find that history is an important investment determinant and conclude that "sunk-cost hysteresis models are empirically relevant."

None of the papers discussed so far tries to recover the parameters that are underlying the firm's dynamic decision and investment problem. Das (1992) investigates the exit decision of cement producers. Using a dynamic programming, she models the decision of cement producers to retire some of their production capacity and estimates the program structurally. In her model, the production, mothballing, or scrapping decision depends on current profits plus the future expected value of the cement production equipment given this period's action. She finds that the probability of retiring cement production equipment is decreasing in the output price, increasing in the oil price (variable costs), and increasing in age. Interestingly, the probability of mothballing is not affected by age, but it is negatively affected by the output price.

Das, Roberts, and Tybout (2000) is another example of a paper that estimates the firm's dynamic investment problem structurally using micro-level data. This study tries to recover the parameters that determine the decision of Columbian chemical producers to enter the export market. They find that the sunk costs of breaking into the export market are substantial. In addition, they are decreasing in firm size, which can be explained by the fact that larger firms are better placed to enter the export market. As Roberts and Tybout (1997), this study finds that history matters. Exporters are very likely to remain exporters, non-exporters are unlikely to enter the export market.

Overall there is considerable empirical evidence suggesting that sunk costs are an important determinant of investment behavior. Researchers use various methods to establish the importance of sunk costs. Testing the theory indirectly, some investigate whether theoretically predicted relations between certain variables hold. Examples of this method are Paddock et al. (1988), Lambson and Jensen (1995; 1998), and Campa (1993). Others investigate the importance of sunk costs considering a firm's observed investment behavior. Examples are Roberts and Tybout (1997), Das (1992), and Das et al. (2000). With either approach, the empirical results are in favor of sunk costs investment models.

2.4 Summary

Theoretical sunk cost models of investment show that sunk costs and uncertainty are very important determinants of investment behavior. For example, the combination of sunk costs and uncertainty over future market conditions can create investment hysteresis. Moreover, sunk costs and uncertainty can be related to entry and exit triggering market conditions, the relation between a firm's current market participation and its investment history, and a firm's market value. Empirical research exploits such theoretical regularities to test sunk costs models of investment and it finds strong evidence for the hypothesis that sunk costs are an important determinant of investment behavior.

Campa (1993), and Lambson and Jensen (1995; 1998) relate sunk cost proxies to other variables and test whether the expected relationships between sunk costs and these variables hold. However, they do not directly test how sunk costs affect a firm's investment behavior. Roberts and Tybout (1997) test for the relevance of sunk cost models using a firm's investment history. However, they do not investigate how different amounts of sunk costs and uncertainty affect investment. An important extension of the existing literature is to analyze investment behavior at the micro-level as in Roberts and Tybout (1997), while allowing for different levels of sunk costs and uncertainty. Using a new data set this is what this dissertation attempts to do. Further extending Roberts and Tybout (1997), this dissertation distinguishes the non-producing firms as mothballed and shutdown.

Chapter 3

OIL PRODUCTION: ENGINEERING AND ECONOMIC FACTORS

This dissertation attempts to study the role of sunk costs and uncertainty over future market conditions in investment decisions. The oil industry is a suitable test case for sunk cost models of investment. The costs of drilling a well are sunk and future oil prices are uncertain. Using aggregate data from Oklahoma and Texas, this chapter presents some evidence that sunk costs are an important determinant of oil well openings and closings.

The most important empirical results of this dissertation are based on a new oil well data set from Oklahoma. As the majority of oil wells in Oklahoma are so called stripper oil wells, I highlight some important features of stripper well production. In addition, I give a brief and general overview of oil well drilling and oil recovery.

3.1 Well Drilling and Plugging

Today most wells are drilled using rotary rigs. The unique feature of this drilling technique is that the drill bit, the tool that cuts and crushes the rock, rotates. The drill bit, operating on the bottom of the borehole, is attached to a long string of steel pipes, the drill string. The drill string is attached to the rotary table on the surface, which is rotated using a motor. As the drilling gets deeper, pipes are added to the drill string. When the bit needs to be changed, well testing is taking place, or repairs to the borehole are necessary and the whole drill string has to be removed from the hole and unscrewed. This is done in reverse again when drilling continues.

Cutting and debris have to be removed from the bottom of the borehole. This is accomplished by flushing drilling mud down the well. The drilling mud is pumped through the center of the drill string and out of the drill bit. From there it flows back to the surface through the space between the drill string and the wall of the borehole flushing out any waste. Other important functions of the drill mud are to keep the drill bit cool and lubricated, coat the sides of the newly drilled well, and control well pressure. The drilling mud is normally a mixture of water or oil and chemicals.

To prevent the hole from collapsing it needs to be cased. This is done with steel pipes cemented into place. As each successive part of the casing is lowered into the borehole through the previous section, the hole becomes smaller in diameter towards the bottom of the well. Depending on the quality of the oil, for example the content of sulfur, expensive specialty steel may be necessary for casing the well. Moreover, in the long-run the steel casing may wear out and require replacement if there is a lot of sand in the oil. During drilling, a blowout preventer is installed on top of the casing, in case a high pressure gas or oil zone is tapped during drilling.

Simple wells consist of one vertical hole. However, it is possible to drill several wells from one location by diverting the angle of each well from the vertical. This is called directional drilling and most often done in off-shore development. An even more complicated drilling technique is horizontal drilling. Here wells are drilled off of the vertical at a certain depth. Horizontal wells are more expensive to drill than non-horizontal wells. Moreover, off-shore wells are much more expensive to drill than on-shore wells. For example, in 1997 the average costs per foot of oil well drilled is $74.23 on-shore and $526.4 off-shore. The average costs per oil well are $374,418 and $5,261,869, on-shore and off-shore respectively, in 1997.[6]

---

[6] See American Petroleum Institute (1999a), Section III, Tables 4 and 5.

There are roughly two types of well drilling: exploratory and development drillings. Exploratory wells are drilled to search for oil, or to gather more information about a newly-discovered oil reservoir. Development wells are drilled for oil extraction [7]in areas that have been proven productive. The number of exploratory wells, as a percentage of total wells drilled in the U.S. in 1997, is 12.74. About thirty percent of the exploratory drillings are successful in 1997. In Oklahoma, a state with old and well developed oil fields, exploratory wells account for 6.6 percent of the number of wells drilled.

At the end of their production lives, wells are plugged and abandoned. Sometimes oil producers decide to temporarily shut-in oil wells (mothball). This is often done when the oil price is very low. State oil and gas regulations normally require mothballed wells to be plugged and abandoned after a set period of time of non-production. Plugging and abandoning is accomplished by filling all producing and water-bearing zones with cement. The well's metal casing is removed and sold if this is profitable.

The costs of drilling an oil well are substantial and mostly sunk. The drilling costs depend on several factors, such as the well depth, the geological environment, and the drilling technique that is used. Not all well drillings are successful and the majority of exploratory wells turn out to be dry. Since wells are plugged with concrete, their re-opening is expensive. These are some of the factors that oil firms have to take into account when contemplating the drilling or plugging of a well.

---

[7] See American Petroleum Institute (1999a), Section III, Table 10.

3.2 Oil Recovery

Under certain conditions geological structures bear oil and develop into oil reservoirs. Oil reservoirs develop when oil accumulates in a formation of porous rock trapped by an impermeable geological structure. Oil reservoirs usually contain oil, gas, and salt water. An oil reservoir may be structured in the following way: the lower part of the porous rock is occupied by water. Above the water is a layer of oil with a gas cap on top. More gas is dissolved in the crude oil. In any case, each oil reservoir has its unique characteristics. The amount of oil that ultimately can be recovered and the appropriate techniques to do so are determined by the reservoir's characteristics. Two very important characteristics are porosity and permeability. The higher the porosity the more oil there is in a given volume of rock and the higher the permeability the better oil can flow through the rock.

Another important reservoir characteristic is its pressure. If the internal pressure on these fluids is larger than the surface pressure, oil is driven spontaneously out of the well. Geologists distinguish three types of natural drive that is in work during the primary recovery stage: water drive, solution gas drive, and gas cap drive. An oil well is water driven if oil is pressed out of the well by water that is captured underneath the oil. In this form of primary recovery, water displaces oil upwards during production and helps to maintain reservoir pressure. In a solution gas driven well it is the gas dissolved in the oil that provides the drive. Solution gas drive works analogous to what happens if a bottle of soda is shaken and opened quickly. As oil is produced, pressure is lost and eventually only gas comes out of the solution. In the third form of primary recovery, the gas cap drive, it is the free gas on top of the oil that forces the oil out of the reservoir. Expanding gas pushes the oil downwards and into the well, which extends into the oil-bearing zone.

With either method of production it is important to control the production rate and pressure decline in order to ensure optimal oil recovery. Potential cumulative oil production is lost if production takes place at a rate too high. Petroleum engineers maximize the cumulative production using a production rule called the efficient rate of production. The efficient rate of production depends on variables such as reservoir size, reservoir pressure, oil viscosity, and rock permeability. In any case, a producing reservoir will eventually reach the point where the natural pressure is too low to push oil to the surface. The worldwide average of oil recoverable by natural drive is about 25 percent of the oil initially in place.[8]

There are several ways of extracting oil beyond the point where the natural drive would fail to work. For example, pressure can be maintained by either injecting water or gas into the reservoir. In this form of secondary recovery, extracted oil is replaced by another fluid. Another type of secondary production is sub-surface pumping. Here oil is actively pumped to the surface. Of course, these artificial methods can be combined. About 96 percent of U.S. oil wells that produce in 1997 use artificial lift methods.[9] Using artificial lifting techniques, about 50 percent of the oil initially in place can be recovered.[10]

Production can be further enhanced by thermal and chemical techniques. Thermal techniques heat up the subsurface oil in order to increase its viscosity and make it flow through the rock more easily. The reservoir temperature can be increased by injecting heat (steam or hot water) or by injecting air into the reservoir and controlled burning of oil. Chemical techniques are used in combination with gas or water injection. For example, water can be thickened so that it replaces the oil more evenly. In any case, enhanced recovery is costly.

---

[8] See Stevenson (1988), p. 156.
[9] See Department of Energy Information Administration (1999), p. 22.
[10] See Stevenson (1988), p. 156.

There are oil wells with natural and artificial flow. Whether oil flows out of the well naturally or artificially depends on the reservoir pressure and other geological variables. Reservoir pressure is lost due to cumulative production and eventually the natural flow ceases to work. This is when more expensive artificial lifting methods are required for oil recovery. Today most oil wells in the U.S. rely on artificial lifting methods with their high variable costs of production.

3.3 Stripper Oil Wells

Oil field production rates generally decrease over time. This is because of depletion and lost pressure. As an oil field ages, the amount of oil produced per well decreases and oil lifting costs increase. Most on-shore oil production in the lower-48 states is relatively old.[11] For example, the currently most productive fields in Oklahoma, Sho-Vel-Tum, and Texas, Wasson, were discovered in 1905 and 1937, respectively.[12] Also, Texas and Oklahoma are the two states with the largest number of stripper wells and the largest amount of stripper well production.[13] In fact, the number of stripper wells is so high that the Department of Energy concludes "Texas, like most of the lower-48 State, is basically a stripper well State."[14] Stripper oil wells are generally defined as wells that produce less than 10 bbls/day. Stripper oil wells are mostly wells in the last stage of their productive life. Generally, oil does not flow out of stripper wells and needs to be pumped to the surface. Because of their high variable costs of production, stripper wells are believed to be only marginally profitable.

---

[11] See Department of Energy Information Administration (2000), Appendix B.
[12] The Department of Energy Information Administration (2000) ranks oil fields by their 1999 production rate.
[13] See Interstate Oil and Gas Compact Commission (1999b).
[14] See Department of Energy Information Administration (2000), p. ix.

Taking a closer look at U.S. oil production in 1997 we see that there were 573,504 producing oil wells in operation and 436,084 oil wells were classified as stripper oil wells.[15] Although these wells produce on average only 2.07 bbls/day in 1998, their total combined production accounted for more than 316 million bbls, or 26 percent of the oil produced in the continental U.S. in the same year.[16] In 1971 the ratio of stripper wells to all other producing wells is 69 percent. As most oil fields have grown older, this ratio increases to 76 percent in 1997. On the whole it is reasonable to conclude that stripper oil wells are an important supplier of crude oil to the U.S. market.

The picture in Oklahoma is similar. In total, there are 88,144 producing oil wells in 1997. In that year, 67,498 oil wells are classified as stripper wells producing, on average, 2.05 bbls of oil per day.[17] There was no well that produced with natural flow in Oklahoma in 1997.[18] Estimates of the average operating cost of oil production for marginal wells in Oklahoma range between \$9.64/bbl and \$15.34/bbl.[19] Especially, in the years 1986 and 1998, when Oklahoma well-head crude oil prices were in the range of \$10 to \$12, marginal oil producers were believed to be under severe economic pressure to abandon wells.

Understanding the investment and production patterns of these small wells is important for several reasons. During periods of low prices the high variable cost stripper wells are in jeopardy of being plugged and abandoned. Because of their high variable costs of production and

---

[15] American Petroleum Institute (1999a), Section III, Tables 13 and 15.
[16] The percentage of 26.6% refers only to continental U.S. oil production. Excluded are Alaska, Florida, and federal offshore, which have no stripper well production. This information is from Interstate Oil and Gas Compact Commission (1999b). Note, that in total the U.S produce 2,354 million bbls of oil in 1997, of which 15% were produced by stripper wells. See American Petroleum Institute (1999a), Section IV, Tables 5 and 6.
[17] See American Petroleum Institute (1999a), Section III, Tables 14 and 15, and Section IV, Table 4.
[18] See American Petroleum Institute (1999a), Section III, Table 14.
[19] See Penn (1996a) for the lower estimate of \$9.64/bbl. For the estimate of \$15.34/bbl see Interstate Oil and Gas Compact Commission (1999a).

the substantial fixed re-opening costs, stripper wells are unlikely to be reopened once plugged and, thus, oil reserves are lost. In order to develop sound policies that prevent the closing of wells during periods of low oil prices and foster the opening of wells during times of high prices, it is important to understand the determinants of well abandonment.

3.4 Entry, Exit, and Oil Production in Oklahoma and Texas

This section presents some evidence that supports the hypothesis that oil production is a sunk cost industry. *Figure 3–1* through *Figure 3–3* show the patterns of entry and exit of oil wells for the states of Oklahoma and Texas, where much of the U.S. stripper well production is concentrated.[20] The plots graph the numbers of entries and exits in the states of Oklahoma and Texas, and the average annual real oil price for the years 1970 to 1997.[21] *Figure 3–1* are graphs of the number of well entries in Oklahoma and Texas and the real price of oil. *Figure 3–2* displays stripper well exits in these two states and the real price of oil. Finally, *Figure 3–3* displays the total number of wells and the total number of stripper wells in Oklahoma and Texas, and the first graph within a row plots data for Oklahoma.

---

[20] Total U.S. stripper well production in was 352.9 million bbs of oil in 1997. Texas and Oklahoma, the largest and second largest stripper oil producers, accounted for 33.4%and 15.7%, respectively, of that. See, American Petroleum Institute (1999a), Section IV, Tables 4 and 5.

[21] The well time series are from the American Petroleum Institute (1999a): Section III, Table 10 (United States Total Wells Reported as Completed by State); Section III, Table 13 (Producing Oil Wells in the United States by State); Section III, Table 15 (United States Stripper Oil Wells by State); and Section III, Table 16 (United States Stripper Oil Wells Abandonments by State). The time-series for stripper wells, Tables 15 and 16, start in 1971. Accordingly the plots for these series start one year later than all other plots.
The average annual real oil price is calculated as the annual average of the monthly spot oil price for West Texas Intermediate, Dollar per Barrel, deflated by the seasonally adjusted producer price index for intermediate energy goods as calculated by the U.S. Department of Labor, Bureau of Labor Statistics. The data source for both time series is the Federal Reserve Economic Database, St. Louis.

Figure 3–1 shows that the number of well completions and the price of oil are highly correlated and exhibit the same pattern over time. In both states the number of well completions is the highest when the oil price is the highest and the number of well completions is lowest when the oil price is low. *Figure 3–2* displays the number of abandoned stripper wells and the price of oil over time. Again, the price of oil is plotted as the solid line. The number of stripper wells abandoned and the oil price are inversely related and we observe the lowest number of exits in the year when the price of oil peaks. The first two panels indicate clearly that entry and exit respond to changes in the oil price, although the exit pattern is less closely related to the price than the entry pattern is. Nevertheless, the price of oil appears to be an important factor in the decision to drill or to abandon an oil well.[22]

The final two graphs in *Figure 3–3* show the total number of producing wells and the number of stripper wells over time in Oklahoma and Texas.[23] The two time series evolve similarly. This indicates that the movements over time in the total number of wells are mainly driven by entries and exits of stripper wells. Moreover, we learn from the last two graphs that the number of wells increases during the time of high oil prices, however, it does not fall back to its initial level when the price falls below the level that triggered the drilling of new wells. The numbers of producing wells and stripper wells in Oklahoma and Texas are always higher after the second oil price shock than before. This is surprising because the real oil prices in 1993, 1994 and 1995 are lower than in the years before the second oil price shocks. However, sunk cost investment models are able to explain this pattern.

---

[22] Note that the number of well completions, *Figure 3–1*, starts to increase in the mid-1970's, before the second oil price shock. This is probably because of increased well profitability due to production and price controls that phased out during that time.

[23] Note the extreme drop in the total number of wells that the time series for Oklahoma displays in 1990. According to American Petroleum Institute (1999a), Section III, Table 14, this is because of a change in the number of oil wells producing with artificial lift. However, after looking at the same series for other U.S. states I conclude that there is a data error in 1990.

This pattern of asymmetric adjustment with respect to price shocks is an example of hysteresis. The pattern is particularly strong when we take into account that oil well production rates generally decline over time and that the variable costs of production increase over time. The plots provide evidence that the number of wells in production is subject to hysteresis and suggests that sunk costs and uncertainty may be important determinants of oil well opening and closing decisions.

In summary, the time-series patterns for the number of wells, entries, and exits in Oklahoma and Texas show the following:

- There is a strong positive relation between the number of new wells drilled and the oil price. The simple correlation between the real oil price and the number of well drillings is .90 and .89 for Oklahoma and Texas, respectively.
- Stripper well exits and the oil price are inversely related. The correlation between the number of stripper well exits and the real oil price is -.60 for Oklahoma and -.53 for Texas.
- The total number of wells does not fall back to its initial level when the price plunges below the entry triggering level. For example, the real oil prices in 1974 ($12.6/bbl) and 1996 ($11.72/bbl) are almost at the same level. Nevertheless, there are 27 percent more oil wells in Oklahoma and 9 percent more oil wells in Texas in 1996 than there are in 1974.
- The time-series evolve similarly in Oklahoma and Texas.

Overall, oil production in Oklahoma appears to be an ideal test case for sunk cost investment models. The small producers in Oklahoma are price takers and strategic considerations among producers do not play a role in the production decision. As needed for testing the theory, there were large fluctuations in oil prices and the costs of drilling an oil well are sunk. Most wells in Oklahoma are only marginally profitable; for many wells the variable costs of production are supposedly between the lowest and highest oil price observed over the period under investigation. If the theory is correct we should only observe a significant increase in the number of well drillings when oil prices are substantially higher than the variable production costs. Similarly, significant well closing should only be observed at oil prices substantially lower than the variable costs of production.

3.5 Summary

Most oil wells in Oklahoma are stripper oil wells with high variable costs of production. In fact, their variable costs of production per barrel of oil are believed to be higher than the price of oil during periods of low prices. Nevertheless, there are not too many oil well abandonments during times of low prices. At the same time, very large numbers of oil well drillings are only observed during times of very high oil prices such as the second oil price shock. This pattern of entry and exit is indicative of hysteresis. Time-series data from Oklahoma and Texas suggests that sunk costs, being able to generate hysteresis, play an important role in the decisions to open and close oil wells.

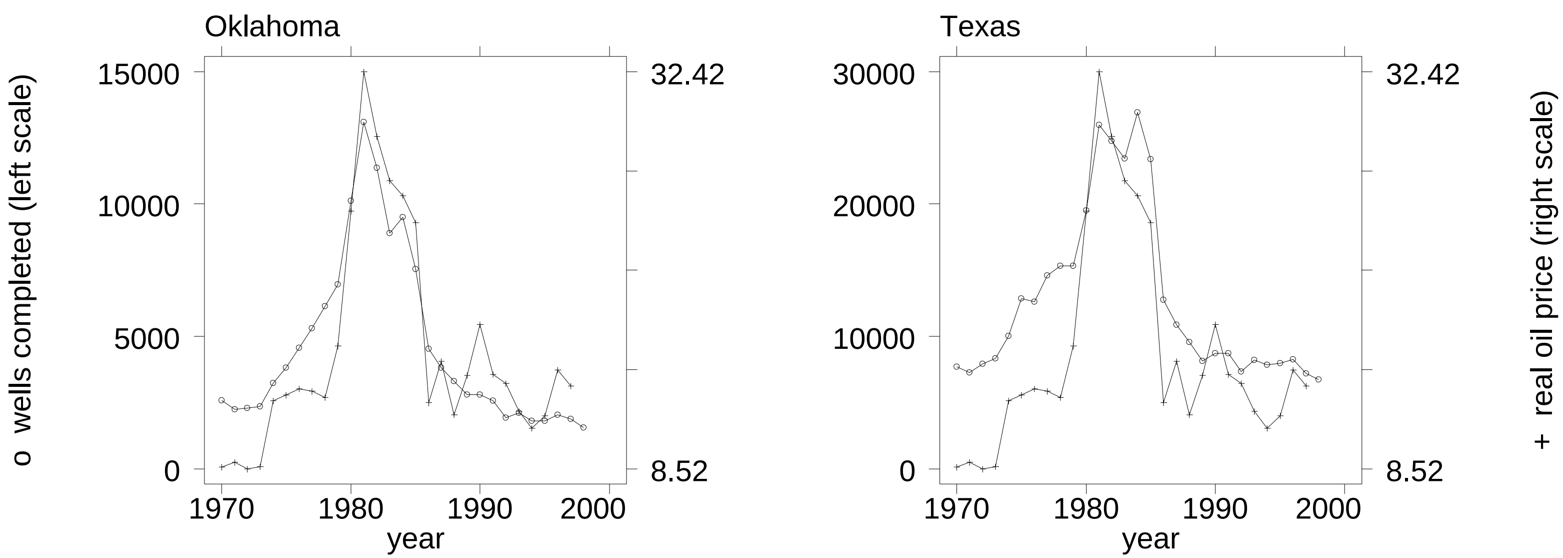

Figure 3–1: Entry of Oil Wells in Oklahoma and Texas, 1970 to 1998

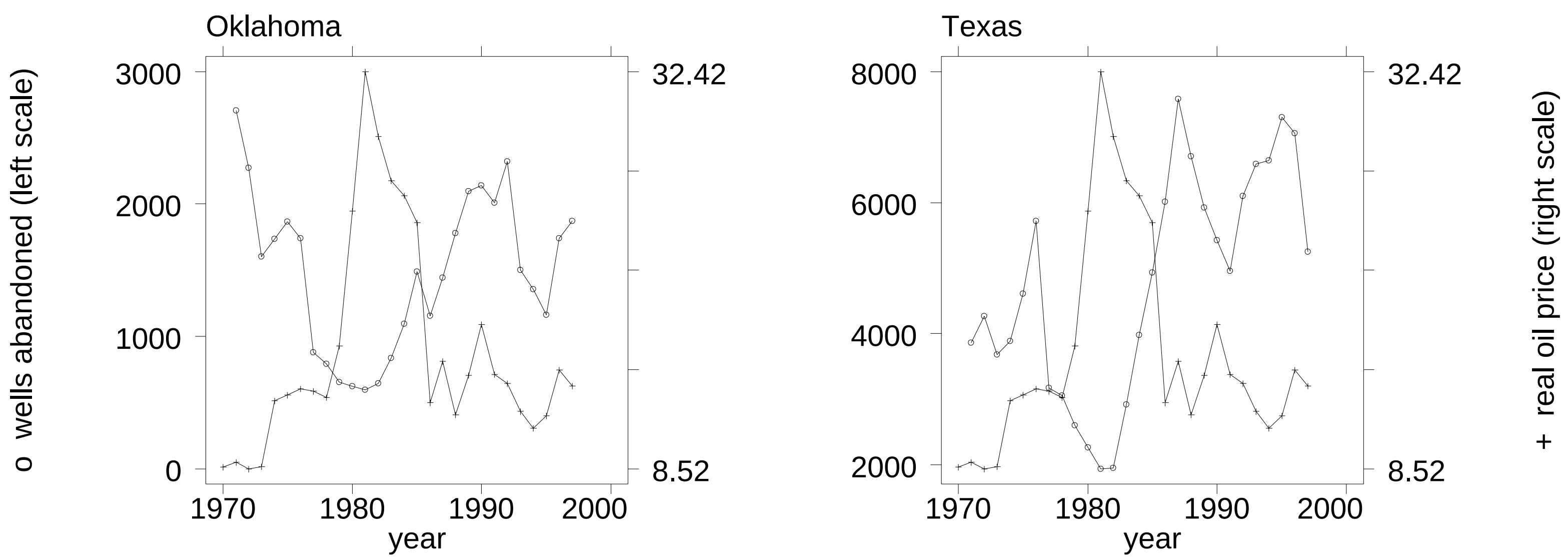

Figure 3–2: Exit of Stripper Wells in Oklahoma and Texas, 1970 to 1998

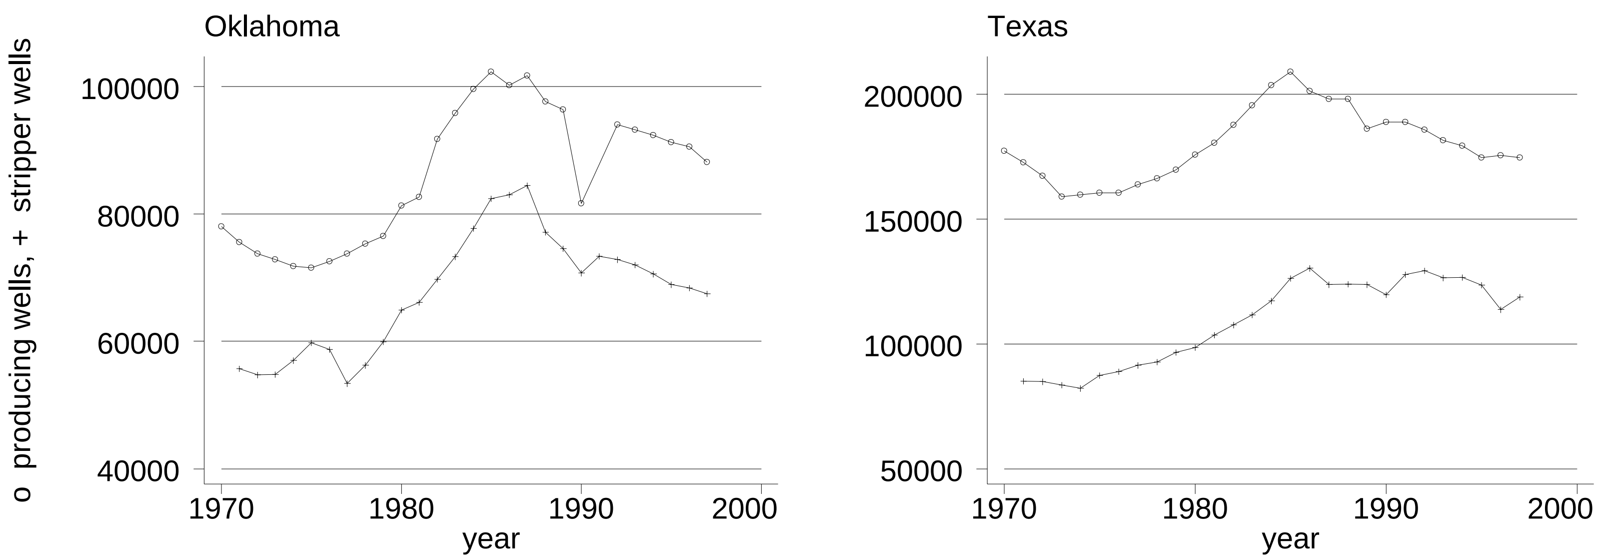

Figure 3–3: Total Number of Producing and Stripper Wells for Oklahoma and Texas, 1970 to 1998

Chapter 4

AN ECONOMIC MODEL OF OIL PRODUCTION

This chapter develops an empirical dynamic, discrete-choice model of oil production, mothballing, and shutdown. The model is built on the assumption that sunk costs are an important determinant in the production decisions. The empirical model will be used to test this hypothesis using micro data on oil production in Oklahoma from 1979 to 1997.

Section 4.1 develops the theoretical model that describes the firm's decision to open, mothball, or shutdown an oil well. Section 4.2 develops the empirical model that will be estimated. Section 4.3 discusses the econometric issues involved in estimating the empirical model.

4.1 The Theoretical Model of Production, Mothballing, and Shutdown

The unit of observation in my data is an oil lease. An oil lease is a geographically confined area for which the lessee has the rights to extract the oil that lies underneath the surface. The lessee has the right to develop the lease without any obligation to drill for an agreed term of time. In general, once production has started the lessee maintains the lease for as long as it is in production and pays royalties to the owner of the land. Oil production on a lease takes place from one or more wells. As discussed in Chapter 3, the opening of a well involves a non-recoverable investment cost that includes the costs of drilling the well and casing it. Existing wells can be shut down. Shutdown is accomplished by plugging the well with concrete. Re-opening of a plugged well is very costly, and it may be optimal to take a well temporarily out of production without plugging. This is called mothballing. There are variable costs incurred when a well is mothballed and there is a legal maximum time length, generally one year, for mothballing. After this time a mothballed well must be plugged.

The data set used in this research includes information on a large number of leases in Oklahoma. The information includes the amount of oil produced from the lease each year and the number of wells drilled or plugged on the lease during the year. This information enables me to define three production states for each lease. A lease is in production if there is at least one active well producing during the year. A lease is mothballed if there is no production on the lease but at least one well on the lease is open. A lease is shut down if all wells on the lease are plugged and abandoned.[24] I will now develop a dynamic programming model that describes the firm's decision to produce from, mothball, or shutdown the lease.[25]

First consider an active, producing lease. If lease *i* is in production in period *t* then $\Pi_{it}(P_t, X_{it}, \theta)$ are its production profits in period *t*. Production profits Π_{it} are a function of market conditions P_t and the characteristics of the lease X_{it}. The vector P_t contains variables such as the price of oil and a measure of uncertainty over future oil prices. X_{it} contains lease specific cost and demand factors. These include the amount of oil produced in prior years, the number of service wells on a lease, and the lease's age. The vector θ contains parameters of the profit function.

---

[24] Although capital adjustment, the decision to increase or decrease the number of wells on a lease, are observed in the data the discrete choice model developed below does not allow for them. The number of capital adjustments in the data is fairly small. Attempts to estimate a model with five production states that include the decisions to increase or decrease the number of wells on a lease were not fruitful.

[25] See Dixit & Pindyck (1994), Chapter 7, for a continuous time version of a related investment problem with the decisions entry/exit and mothball/produce.

Switching the state the lease occupies entails sunk costs. These are non-recoverable costs that are paid only when the lease switches among the three states. They are summarized in *Table 4–1*.

Table 4–1: Sunk Costs of Changing the Production State

Sunk Costs of Changing the Production State				
	Lease Status in Period t			
		Production	Mothballing	Shutdown
Lease Status in Period $t\text{-}1$	Production	0	M_i	M_i+S_i
	Mothballing	RM_i	0	S_i
	Shutdown	RM_i+RS_i	RS_i	0

As shown in the first row, the costs of switching from production to mothballing are M_i. A producer can skip the mothballing stage and move directly to shutdown at costs M_i+S_i. The second row summarizes the switching costs for leases that were mothballed in period $t\text{-}1$. A mothballed lease can switch to production or to permanent shutdown. A producer who has mothballed a lease can re-enter production by paying the re-entry cost RM_i. The costs of switching from mothballing to shutdown are S_i. Finally, as shown in the last row, a producer can skip the mothballing stage and move directly between production and shutdown. The costs of re-entering production are RM_i+RS_i, and the costs of moving from shutdown to mothballing are RS_i. The non-recoverable switching costs are treated as lease specific and may vary across leases. The empirical model for instance, uses the average well depth on a lease to allow sunk costs to vary with lease characteristics.

The actual per period profit s realized by the lease in each of the three possible states in year t depend on the market and lease conditions (P_t, X_{it}) in year t and the production status of the lease in period $t-1$. Given the sunk costs, the per period operating profit from an active lease $\Pi_{it}(P_t, X_{it}, \theta)$, and a per period mothballing costs vm_i, the actual profit earned on a lease can be described by:

$$
\begin{aligned}
\text{produce} &: \Pi_{it}(P_t, X_{it}, \theta) - m_{it-1} RM_i - s_{it-1}(RM_i + RS_i) + \varepsilon^*_{i1t} \\
\text{mothball} &: -vm_i - (1 - m_{it-1}) M_i - s_{it-1}(RS_i - M_i) + \varepsilon^*_{i2t} \qquad (4.1) \\
\text{shutdown} &: -(1 - s_{it-1})(M_i + S_i) + m_{it-1} M_i + \varepsilon^*_{i3t}.
\end{aligned}
$$

The indicator variable m_{it-1} takes the value 1 if the lease is mothballed in period $t-1$ and zero otherwise. The indicator variable s_{it-1} takes the value 1 if the lease is shutdown in period $t-1$ and zero otherwise. In addition, the specification in equation 4.1 assumes that there are random profit shocks denoted ε^*_{i1t}, ε^*_{i2t}, and ε^*_{i3t}.

Equation 4.1 is straightforward. In the producing state the lease generates production profits Π_{it}. In addition, if the lease did not produce in the previous period, there are re-entry costs RM_i if returning from mothballing, and RM_i+RS_i if returning from shutdown. The last two equations describe current profits from non-production. If the lease is mothballed the lessee must pay variable costs vm_{it}. If the lease switched from production to mothballing, the additional cost M_i is incurred. The cost of switching from shutdown to mothballing is (RS_i-M_i). The last line in 4.1 describes shutdown costs. The are no variable costs or profits from shutdown. Switches from the state production to shutdown cost M_i+S_i and switches from mothballing to shutdown cost S_i.

The current profit equations can be nested as:

$$\begin{aligned}
&R_{it}(P_t, X_{it}, Y_{it}, \theta) = \\
&(1-m_{it})(1-s_{it})\left\{\Pi_{it}(P_t, X_{it}, \theta) - m_{it-1}RM_i - s_{it-1}(RM_i + RS_i) + \varepsilon^*_{i1t}\right\} \\
&- m_{it}\left\{\nu m_i + (1-m_{it-1})M_i - s_{it-1}(RS_i - M_i) + \varepsilon^*_{i2t}\right\} \\
&- s_{it}\left\{(1-s_{it-1})(S_i + M_i) - m_{it-1}M_i + \varepsilon^*_{i3t}\right\}.
\end{aligned} \qquad (4.2)$$

Define $Y_{it} \equiv (m_{it}, s_{it})'$. At most one of the dummy variables can take the value 1 for firm *i* in period *t*. The vector Y_{it} takes the values

$Y_{it} = \left\{\begin{pmatrix}0\\0\end{pmatrix}, \begin{pmatrix}1\\0\end{pmatrix}, or \begin{pmatrix}0\\1\end{pmatrix}\right\}$, denoting the production, mothballing, or shutdown state respectively.

Each period the lease owner is assumed to choose the sequence of decision rules

$Y^{(+)}_{it} = \{Y_{it} = f_t(P_t, X_{it}, \theta)\}^{\infty}_{t=0}$ that maximizes the expected present value of payoffs in period *t*.

The maximized expected present value is represented by the lease's value function

$$V_{it}(\Omega_{it}) = \max_{Y^{(+)}_{it}} E_t\left(\sum_{l=t}^{\infty} \beta^{l-t} R_{il}(P_l, X_{il}, Y_{il}, \varepsilon^*_l, \theta) \middle| \Omega_{it}\right) \qquad (4.3)$$

where β it the one period discount rate and expectations are conditioned on the lease specific information set Ω_{it}. Using Bellman's equation the value function in equation 4.3 can be rewritten as

$$V_{it}(\Omega_{it}) = \max_{Y_{it}}\left(R_{it}(P_t, X_{it}, Y_{it}, \varepsilon^*_{it}, \theta) + \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it}\right]\right), \qquad (4.4)$$

which implies that the lease value in period *t* is equal to current profit plus the discounted expected period *t+1* lease value given that the profit and value maximizing state is chosen in period *t*.

The values of the Bellman equation corresponding to the three possible choices of Y_{it} are

$$\begin{aligned}
\text{produce}: Z_{it}^{*1} &= \Pi_{it}(P_t, X_{it}, \theta) - m_{it-1}RM_i - s_{it-1}(RM_i + RS_i) \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (0,0)'\right] + \varepsilon_{i1t}^{*} \\
\\
\text{mothball}: Z_{it}^{*2} &= -vm_i - (1 - m_{it-1})M_i - s_{it-1}(RS_i - M_i) \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (1,0)'\right] + \varepsilon_{i2t}^{*} \\
\\
\text{shutdown}: Z_{it}^{*3} &= -(1 - s_{it})(S_i + M_i) + (1 - s_{it-1})m_{it-1}M_i \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (0,1)'\right] + \varepsilon_{i3t}^{*} .
\end{aligned} \qquad (4.5)$$

A lease owner will chose the action with the largest current expected value. Therefore choice j is observed in period t if $Z_{it}^{*j} > Z_{it}^{*k} \ \forall\, k \neq j$. Note that because of the variable mothballing costs it is never optimal to switch from shutdown to mothballing. Thus, the shutdown dummy variable in the mothballing equation is always zero. After adjusting for this equation 4.5 becomes

$$\begin{aligned}
\text{produce}: Z_{it}^{*1} &= \Pi_{it}(P_t, X_{it}, \theta) - m_{it-1}RM_i - s_{it-1}(RM_i + RS_i) \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (0,0)'\right] + \varepsilon_{i1t}^{*} \\
\\
\text{mothball}: Z_{it}^{*2} &= -vm_i - (1 - m_{it-1})M_i \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (1,0)'\right] + \varepsilon_{i2t}^{*} \\
\\
\text{shutdown}: Z_{it}^{*3} &= -(1 - s_{it})(S_i + M_i) + (1 - s_{it-1})m_{it-1}M_i \\
&+ \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it} = (0,1)'\right] + \varepsilon_{i3t}^{*} .
\end{aligned} \qquad (4.6)$$

Next, note that subtracting one choice from the other choices would not change anything in the relative profit ordering. Only the actual choice in period t is observed and this set of equations is empirically not identified. To generate a profit relation that can be estimated, I subtract the first category, “produce”, from the other choices.

To facilitate the algebraic expression define

$$\Delta V^{e}_{i,t+1}(m_{it}=1) \equiv \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it}=(1,0)\right] - \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it}=(0,0)'\right],$$

$$\Delta V^{e}_{i,t+1}(s_{it}=1) \equiv \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it}=(0,1)'\right] - \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1}) \mid Y_{it}=(0,0)'\right].$$

This is the difference in the discounted expected lease values in period t+1 of the mothballing and the production state, and the shutdown and the production state. Similarly define the differences in random profit shocks as $\varepsilon_{i2t} \equiv \varepsilon^{*}_{i2t} - \varepsilon^{*}_{i1t}, \varepsilon_{i3t} \equiv \varepsilon^{*}_{i3t} - \varepsilon^{*}_{i1t}$.

Rearranging terms, recognizing that m_{it-1}=1 implies that (1-s_{it-1})=1, and normalizing the system the by the first choice gives:

$$\text{produce}: Z^{1}_{it} = 0$$

$$\begin{aligned}\text{mothball}: Z^{2}_{it} = &-vm_i - M_i - \Pi_{it} + \beta\Delta V^{e}_{i,t+1}(m_{it}=1)\\ &+ m_{it-1}(M_i + RM_i) + \varepsilon_{i2t}\end{aligned} \qquad (4.7)$$

$$\begin{aligned}\text{shutdown}: Z^{3}_{it} = &-(S_i + M_i) - \Pi_{it} + \beta\Delta V^{e}_{i,t+1}(s_{it}=1)\\ &+ m_{it-1}(M_i + RM_i) + s_{it-1}(S_i + M_i + RM_i + RS_i) + \varepsilon_{i3t}.\end{aligned}$$

This system of equations is in the form of a multinomial discrete choice model and is the basis for the empirical estimation in this dissertation.

Similar to Roberts and Tybout's (1997) binary investment model, one can test for sunk costs by testing for state dependence. In the presence of sunk costs, namely M_i>0, S$_i$>0, RM_i>0, and RS_i>0, the lagged production state, indicated by m_{it-1} and s_{it-1}, matters for the current choice of production state. By estimating the above set of equations and testing whether coefficients on the lagged state dummy variables are significant and positive, I can test whether sunk costs matter.

What if there are no sunk costs involved in moving from one state to another, that is M_i=S_i=RM_i=RS_i= 0?[26] Switching between states is costless and instantaneous. In this case changes in expected future values vanish $(\Delta V_{i,t+1}^{e}=0)$ because the current production state does not matter for future decisions. The investment decision collapses to a purely static one because the firm can switch between states instantaneously and free of any costs. If there are no sunk costs, the firm will produce if the current profits Π_{it} are positive and shutdown if they are negative. There will be no mothballing because of the variable mothballing costs vm_i. On the other hand, the larger are the non-recoverable switching costs, the more important is the lagged production status for the current period production choice.

Because changes in the parameters governing the oil price evolution, and thus changes in expected future oil prices, have implications for investment decisions only through the $\Delta V_{i,t+1}^{e}$ terms, changes in these parameters matter only if sunk costs are non zero. In particular, changes in oil price volatility will affect the option value of investment and the investment decision only if sunk costs matter. Thus, including oil price volatility in the estimation and testing for its significance is an indirect way to test whether sunk costs matter.

A general theoretical finding is that the investment option value increases with uncertainty. What is the value of the investment opportunity for choice j=2 and 3? Here the value of the opportunity depends on the lagged production status of the lease. A non-producing lease has the opportunity to return to production and vice versa. Consequently, the probability of being mothballed or shut down is increasing in uncertainty if the lease did not produce in the previous period and decreasing in uncertainty if the lease was in the producing state in the previous period.

---

[26] Note, with this sunk cost specification the mothballing state will not be used at all.

4.2 The Empirical Model

This section discusses the estimation strategy for the model developed in section 4.1 and derives an estimation. The estimation equation is based on the model outlined in equation 4.7, which identifies the conditions that determine the state of a lease.

In the choice model of equation 4.7, the lagged choice indicators s_{it-1} and m_{it-1} enter the current choices directly. The variables in P_t and X_{it} enter the model through the current period production profit function Π_{it}, and the expected, next period value function $V^e_{i,t+1}$. Instead of assuming specific functional forms for the production profit and the value functions, it is assumed that they are sufficiently approximated with a reduced from expression. Π_{it}, and $V^e_{i,t+1}$ are treated as latent variables that differ across leases and time because of variations in lease characteristics and macro-variables. I specify that variations in the latent profit and future expected lease value differences, across leases and time, arise from changes in the macro level variables $\boldsymbol{P_t}$, the lease characteristics $\boldsymbol{X_{it}}$, and random noise ε_{ijt} in a linear fashion. Additionally, there are unobservable sunk costs and variable costs terms that are not attached to m_{it-1} and s_{it-1}. In the empirical model, these terms are captured in the intercepts. With these assumptions the latent profits and lease values in Z^2_{it} and Z^3_{it} in equation 4.7 can be expressed as:

$$\begin{aligned}
\text{in } Z^2_{it} &: -vm_i - M_i - \Pi_{it} + \beta\Delta V^e_{i,t+1}(m_{it}=1) \\
&= \alpha_2 + \boldsymbol{\beta}_2\boldsymbol{P}_t + \boldsymbol{\delta}_2\boldsymbol{X}_{\boldsymbol{it}} + \varepsilon_{i2t} \\
\text{in } Z^3_{it} &: -(S_i + M_i) - \Pi_{it} + \beta\Delta V^e_{i,t+1}(s_{it}=1) \\
&= \alpha_3 + \boldsymbol{\beta}_3\boldsymbol{P}_t + \boldsymbol{\delta}_3\boldsymbol{X}_{\boldsymbol{it}} + \varepsilon_{i3t}\ .
\end{aligned} \qquad (4.8)$$

The remaining terms in equation 4.7 are the lagged discrete production indicators s_{it-1} and m_{it-1} multiplied by entry or exit costs. Although these costs are not observable, it is possible to test for their presence by estimating coefficients on s_{it-1} and m_{it-1}. The specification of entry and exit costs can be improved upon by recognizing that they vary with some lease characteristics. For example, the average depth of the wells on a lease can serve as a proxy for sunk costs. Combining the lagged dummy variables with this proxy allows sunk costs to be different across leases. Another proxy for sunk costs is the number of wells on a lease. The more wells there are on a lease, the more costly it is to shut down. Therefore, the number of wells on a lease is another variable that can be combined with the lagged state dummy variable. Since the hysteresis band increases in sunk costs the coefficient on all these variables are expected to be positive.

For estimating the coefficients on the sunk costs proxies, an intercept term, the average well depth, and the number of wells on a lease are organized in the vector H_{it} and interacted with the lagged production status dummy s_{it-1} or m_{it-1}. The estimation equation that represents the pay-offs to each choice j is:

$$\text{produce}: Z_{it}^{1} = 0$$

$$\text{mothball}: Z_{it}^{2} = \alpha_2 + \boldsymbol{\beta}_2 \boldsymbol{P}_t + \boldsymbol{\delta}_2 \boldsymbol{X}_{\boldsymbol{it}} + \boldsymbol{\gamma}_2 m_{it-1} \boldsymbol{H}_{\boldsymbol{it}} + \varepsilon_{i2t} \qquad (4.9)$$

$$\text{shutdown}: Z_{it}^{3} = \alpha_3 + \boldsymbol{\beta}_3 \boldsymbol{P}_{\boldsymbol{t}} + \boldsymbol{\delta}_3 \boldsymbol{X}_{\boldsymbol{it}} + \boldsymbol{\gamma}_3 m_{it-1} \boldsymbol{H}_{\boldsymbol{it}} + \boldsymbol{\tau}_3 s_{it-1} \boldsymbol{H}_{\boldsymbol{it}} + \varepsilon_{i3t}.$$

Choice j is observed if $Z_{it}^{j} > Z_{it}^{k}\ \forall\ j \neq k$, j,k=1,2,3.

Finally, for estimating equation 4.9 it is necessary to make assumption regarding the random noise ε_{jit}. For example, ε_{jit} can be logit or normal. The random noise term, along with other econometric issues, is discussed in section 4.3.

4.2.1 A Special Case: Production and Non Production States

The models developed in section 4.1 allow the firm to choose among three different activities: production, mothballing, and shutdown. However, many times it is hard to distinguish between shutdown and mothballing. For example, a data set may just contain production data. In this case, it is impossible to tell whether zero production is due to mothballing or shutdown. For this data problem, this section develops a simpler model that distinguishes only two states: production and no-production. How does the model change if only the two states production and no-production are observed? Equation 4.2 expresses the current profits in the model with all three choices. This expression can be manipulated by assuming that some cost terms are zero. If it is not possible to empirically distinguish between mothballing and shutdown it is sensible to assume that these two states are the same, e.g. that there is no difference between the costs of mothballing and shutdown, specifically vm_{it}=0, S_{it}=0, and RS_{it}=0. Moreover, it is assumed that the random shocks for mothballing and shutdown are the same. Rearranging terms the static current profits are now:

$$\begin{aligned} R_{it}(\bullet) &= (1-m_{it})(1-s_{it})\left\{\Pi_{it}(P_t, X_{it}) - (m_{it-1}+s_{it-1})RM_i + \varepsilon^*_{i1t}\right\} \\ &\quad - (m_{it}+s_{it})\left\{(M_i(1-m_{it-1}-s_{it-1}) + \varepsilon^*_{i2t})\right\} \end{aligned} \qquad (4.10)$$

In this case M_i and RM_i represent the sunk costs of moving out of production and into production, respectively. Define indicator for no-production np_{it}=1 if m_{it}=1 or s_{it}=1 and zero otherwise. Equation 4.10 can be re-written as:

$$\begin{aligned} R_{it}(\bullet) &= (1-np_{it})\left\{\Pi_{it}(P_t, X_{it}) - np_{it-1} * RM_i + \varepsilon^*_{i1t}\right\} \\ &\quad - np_{it}\left\{M_i * (1-np_i) + \varepsilon^*_{i2t}\right\} \end{aligned} \qquad (4.11)$$

Using the same logic as before, the lease is producing oil if the current and expected future profits from production are larger than the current and expected future profits from not producing. The oil lease's current values for the choices production and no-production are

$$\begin{aligned} &\text{production}: Z_{it}^{1} \equiv \Pi_{it}(\,\bullet\,) - np_{it-1} \times RM_i + \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1})|np_{it} = 0\right] + \varepsilon_{i1t}^{*} \\ &\text{no - production}: Z_{it}^{2} \equiv -(1 - np_{it-1})M_i + \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1})|np_{it} = 1\right] + \varepsilon_{i2t}^{*}\,. \end{aligned} \qquad (4.12)$$

Define the latent variable representing the expected current value increment of not producing today as:

$$NoP_{it}^{*} \equiv -\,\Pi_{it}(\,\bullet\,) + \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1})|np_{it} = 1\right] - \beta E_t\left[V_{i,t+1}(\Omega_{i,t+1})|np_{it} = 0\right]. \qquad (4.13)$$

Define $\varepsilon_{it} \equiv \varepsilon_{i2t}^{*} - \varepsilon_{i1t}^{*}$. Normalizing the set of equations by the choice production the decision not to produce in period *t* can be summarized as:

$$np_{it} = \begin{cases} 1 \text{ if } & NoP_{it}^{*} - M_i + np_{it-1}(M_i + RM_i) + \varepsilon_{it} \geq 0 \\ \\ 0 \text{ otherwise.} \end{cases} \qquad (4.14)$$

A lease does not produce if the current profits and the discounted, expected next period lease value from being in the no-production state are larger than in the production state. This is adjusted for the sunk costs that are occurred when switching between states. The larger the sunk costs are the more likely it is that the expression in the first row of equation 4.14 is larger than zero.

Interpreting latent profits and entry and exit costs analogous to section 4.2 this can be estimated with the same set of variables as the model with three choices. The empirical binary choice equation is

$$np_{it} = \begin{cases} 1 \text{ if } \alpha + \boldsymbol{\beta P}_t + \boldsymbol{\delta X_{it}} + \gamma np_{it-1}\boldsymbol{H}_{it} + \varepsilon_{it} \geq 0 \\ \\ 0 \text{ otherwises.} \end{cases} \qquad (4.15)$$

As in the model with three choices, the coefficients of interest are γ. Again, I expect these coefficients to be positive.

This model does not differentiate between mothballing and shutdown. Although some facets of the investment and production decision are lost, estimating the simpler binary problem model has several advantages. In many data sets it might not be possible to distinguish between mothballing and shutdown. Also, apart from data availability, in many cases it is much easier to estimate a binary choice model than a multinomial choice model.

4.3 Econometric Issues

Section 4.3.1 discusses a desirable error structure for the two and three-choice empirical model developed in section 4.2. Section 4.3.2 discusses a feasible limited dependent variable (LDV) estimator that allows for this error structure.

4.3.1 Error Structure

Testing for sunk costs is done by testing for persistence in the production status of the lease, i.e. state dependence. To estimate the importance of sunk costs correctly, it is crucial to control for other potential sources of persistence in the production status. State dependence potentially arises from current and future expected profit levels. Including the vector of observable lease characteristics X_{it} in the estimation, controls for much of this dependence. Nevertheless, most likely there are some important characteristics that are unobserved, which can induce serial correlation in the error terms ε_{jit}. Furthermore, it is possible that errors are correlated among choices. For instance, a positive profit shock for the production state probably has a negative effect on the values of the mothballing and shutdown states. Therefore, the model is specified with an error structure that allows for a serially-correlated choice component and a random effect lease component.

An error structure that allows for serial correlation and a random effects component in the normalized production, mothballing, and shutdown model is

$$\begin{pmatrix} \varepsilon_{i1t} \\ \varepsilon_{i2t} \end{pmatrix} = \sqrt{1-\lambda} \times \begin{pmatrix} \omega_{i1t} \\ \omega_{i2t} \end{pmatrix} + \sqrt{\lambda} \times \begin{bmatrix} 1 & 0 \\ \psi_{12} & \psi_{22} \end{bmatrix} \begin{pmatrix} \nu_{i1} \\ \nu_{i2} \end{pmatrix}, \text{ where}$$

$$\begin{pmatrix} \omega_{i1t} \\ \omega_{i2t} \end{pmatrix} = \begin{bmatrix} \rho_1 & 0 \\ 0 & \rho_2 \end{bmatrix} \begin{pmatrix} \omega_{i1t-1} \\ \omega_{i2t-1} \end{pmatrix} + \begin{bmatrix} 1 & 0 \\ \phi_{12} & \phi_{22} \end{bmatrix} \begin{pmatrix} \eta_{i1t} \\ \eta_{i2t} \end{pmatrix}. \qquad (4.16)$$

The regression errors $\boldsymbol{\varepsilon}_{it}$ consist of the transitory component $\boldsymbol{\omega}_{it}$ and a permanent lease-specific component $\boldsymbol{\nu}_i$. The transitory component can be correlated across choices (ϕ_{12}) and over time (ρ). The permanent component will also be correlated across choices (ψ_{12}). The weight $\lambda \in [0,1]$ determines the contribution of each component to the regression error $\boldsymbol{\varepsilon}_{it}$. If λ=0 than there is no permanent component. The transitory component is modeled as serially correlated (ρ_1, $\rho_2 \neq 0$) and correlated across choices ($\phi_{12} \neq 0$). The lease-specific permanent component is correlated across choices ($\psi_{12} \neq 0$) and does not vary over time.[27] There are several useful special cases. For example, there is no serial correlation in the transitory component if $\rho_1 = \rho_2 = 0$. Profits and lease values are not correlated if $\phi_{12} = \psi_{12} = 0$.

Assuming that $\boldsymbol{\nu}_i$ and $\boldsymbol{\eta}_{it}$ are normally distributed random variables, this error structure generates a multinomial probit model. Because of the serially correlated errors and the fact that there are more than two choices in each period, estimating this probit model is non-trivial.

[27] For reasons of identification, the variance of the first normalized choice, and thus the parameter in the most upper, left corner of both covariance matrices, is normalized to one. Note, in the binary case these two matrices collapse to a scalar, which is equal to one.

This error structure can be used for the binary production model with the choices production and no-production, too. For the binary case, the regression errors are specified as

$$\varepsilon_{it} = \sqrt{1-\lambda} \times \omega_{it} + \sqrt{\lambda} \times \upsilon_i \text{ , where } \omega_{it} = \rho\omega_{it-1} + \eta_{it} . \qquad (4.17)$$

Again, ρ determines serial correlation in the transitory component and λ weights the permanent component relative to the transitory component. λ=0 and ρ=0 implies the standard probit model with independent observations.

4.3.2 A Multinomial Probit Estimator

Keane (1994) develops a feasible simulation estimator for probit models on panel data that allows for regression errors as the ones specified in section 4.3.1. Generally, the maximum likelihood estimation of panel LDV models requires the evaluation of hard to compute multivariate integrals. Keane's panel probit estimator avoids numerically evaluating these integrals by using simulated choice probabilities. Keane's LDV panel estimator is based on McFadden's (1989) method of simulated moments (MSM). Below is a brief summary of the estimator and the algorithm used for simulating the choice probabilities.

Let d_{ij} be an indicator that individual i chose j, Let p_{ij} be the probability of the event d_{ij}=1, which is a function of unknown parameters θ. The method of moment estimator of $\hat{\theta}$ is formed by solving the moment conditions:

$$\frac{\partial L(\theta)}{\partial \theta} = \sum_{i=1}^{N} \sum_{j=1}^{J} w_{ij} \left(d_{ij} - \hat{p}_{ij} \right) = 0. \qquad (4.18)$$

Using the optimal weighting scheme,

$$w_{ij} = \frac{\partial \ln(\hat{p}_{ij})}{\partial \hat{\theta}} = \frac{\partial \hat{p}_{ij} \big/ \partial \hat{\theta}}{\hat{p}_{ij}} , \qquad (4.19)$$

this estimator is asymptotically equivalent to the ML estimator.

McFadden's MSM estimator simulates the probabilities in 4.18. This method gives consistent estimates of θ. Moreover, the weights are also functions of the simulated probabilities. They do not need to be exact and updated at every iteration. An advantage of this method is that the simulated probabilities enter the maximizing conditions linearly. Because of this, simulation errors tend to cancel out over individuals *i*.

This method can be directly applied to panel data if d_{ij} represents *i*'s sequence of choices over all time periods, t=1,…,T. However, with J choices each period and T time periods there are J^T choice sequences to choose from. For applying MSM directly to the panel case the probability of each sequence needs to be simulated. This estimator quickly becomes impractical when T becomes large. Keane (1994) develops a MSM estimator that is feasible with large J^T. He decomposes the probability of a choice sequence into the product of conditional probabilities. Equations 4.18 and 4.19 can be rewritten as

$$\sum_{i=1}^{N}\sum_{j=1}^{J}\sum_{t=1}^{T_j} w_{ijt}\left(d_{ijt} - \hat{p}_{ijt}^{*}\right), \qquad (4.20)$$

and

$$w_{ij} = \frac{\partial \hat{p}_{ij}^{*} / \partial \hat{\theta}}{\hat{p}_{ij}^{*}}, \qquad (4.21)$$

where $\hat{p}_{ijt}^{*}$ is the conditional probability of making choice j in period t:

$$p_{ijt}^{*} = \text{prob}\left(d_{ijt} = 1 \middle| D_{it-1}\right). \qquad (4.22)$$

The choice history is summarized in D_{it-1} and the transition probability p_{ijt}^{*} is conditioned on all choices made by *i* in periods previous to *t*. Note that the number of simulated probabilities for individual *i* is now $J*T$.

Simulating the conditional transition probabilities, though, is not trivial. Several simulation algorithms have been developed for simulating choice probabilities in LDV models. An efficient and unbiased method for simulating the transition probabilities is the Geweke-Hajivassiliou-Keane (GHK) simulation algorithm. The GHK simulation algorithm draws the choice probabilities in period *t* from an importance sampling distribution, which conditions on period *t-1*. As mentioned earlier, simulating the probability of a whole choice sequence becomes quickly infeasible when *T* becomes large. However, note that the probability of a particular choice sequence can be decomposed into the product of conditional choice probabilities. For example, the probability of the sequence of choices from t=1 to t=k can be rewritten as the probability of the actual choice in period k conditional on the choices made up to k-1, multiplied by the probability of the actual choice in period k-1 conditional on the choices made up to k-2, and so on. The algorithm for simulating these probabilities works recursively starting in the first period.

The GHK probability estimator works along the following sequential procedure: (1) Draw, from the truncated univariate normal distribution, an error term such that the observed choice j is made in t=1. (2) Retain the value drawn in t=1. Conditional on the acceptable draw for the first period, draw for the second period such that the observed choice j is made. (3) Continue with conditional draws as above until there is a sequence of errors for periods 1 to t-1 that generates the sequential choices that from period 1 to t-1. (4) The GHK simulated probability is then the product of the conditional probabilities.

With the GHK simulator the choice probability, i.e. transition probability, that i chooses j in period t, p^*_{ijt}, is then constructed as a ratio of GHK probabilities

$$\hat{P}_{GHK}\left(d_{ijt} \middle| D_{it-1}\right) = \hat{P}_{GHK}(d_{i1},...,d_{it-1},d_{ijt}) \Big/ \hat{P}_{GHK}(d_{i1},...,d_{it-1}). \qquad (4.23)$$

Studies of alternative probability simulators find that the GHK simulator is the most accurate of the ones under consideration (see Keane (1993) for an overview).

A alternative, readily available and easier to implement estimator for the three-choice model in section 4.2 is the multinomial logit estimator. Unlike the multinomial probit estimator, multinomial logit has a convenient closed-form solution and numerical integration is avoided. However, the pattern of predicted probabilities generated by the multinomial logit estimator is restricted and seems unrealistic in many circumstances. The multinomial logit estimator implies that the choice alternatives enter each choices probability equally and that choice alternatives are equally substitutable. This restriction is known as independence of irrelevant alternatives.

Multinomial probit with non-independent errors is capable of generating more realistic choice patters. However, it comes at the cost of numerically difficult integrals. The MSM estimator is one of the methods that have been developed to estimate multinomial probit with a relatively unrestricted covariance matrix and it has several advantages over other multinomial estimators. Another simulation estimator that avoids numerical integration of the choice probabilities is to use simulated probabilities in the likelihood function itself. This approach is known as simulated maximum likelihood (SML). However, attempts of using SML to estimate multinomial probit models with serially correlated regression errors rendered unsatisfactory results. Keane (1994) finds "that SML produces severely biased estimates of the serial correlation structure, while the MSM estimator does not." McFadden (1989) recognizes that the simulated probabilities have to enter the conditions that define the estimator linearly for the estimation strategy to be viable when the number of simulations is small. Unlike MSM, where the simulated probabilities enter the moment conditions linearly (see equation 4.20), the SML estimator violates this condition. For an overview of simulation estimation methods for LDP models see Hajivassiliou (1993).

4.3.3 Initial Conditions

Serial correlation in profit shocks is not the only econometric problem that can lead to biased estimates of coefficients on the lagged state indicators. In the data, a lease's production and investment history is observed for the years 1979 through 1998. However, the production and investment choice in the pre-sample period cannot be treated as exogenous. Because of state dependence and potential serial correlation in the error terms, a lease's state in 1979 is affected by its prior production and investment history. Thus, the condition the lease is in initially matters for the observed history.

Several studies discuss the initial conditions problem for nonlinear models in a parametric framework and propose different solutions. Heckman (1981), for example, suggests approximating the conditional distribution of the initial condition. He proposes to use an approximate representation for the choice in the pre-sample period and to allow for correlation between the disturbances from this pre-sample expression and the disturbances from the sample expression. It can be computationally difficult to estimate this approximation.

Wooldridge (2000) develops another solution to the initial conditions problem. For a general class of nonlinear models with unobserved heterogeneity, he suggests estimating the econometric model conditioned on the initial values and exogenous explanatory variables. The density of the unobserved effect is conditioned on the initial state or choice. This leads to a very simple procedure in the probit case when the correctly-specified density includes only one lag of the dependent variable: add the initial choice observed in the data in year t=1 to the regressors and estimate the probit model with the remaining data for year t=2 to T.

In the empirical model I will use Keane's multinomial probit estimator, together with Wooldridge's initial condition correction to estimate the dynamic discrete choice model developed in this chapter.

4.4 Summary

This chapter develops two empirical models of oil production. (1) A model in which leases are producing, mothballed, or shut down. This three-choice model allows for transitions between all three states and can be estimated using equation 4.9. In this dynamic choice model, the current production status depends on the lagged production status if the costs of switching the production are non-recoverable. (2) The second model allows only for the states production and no-production. In many instances it is empirically not possible to distinguish between mothballing and shutdown. This model covers this special case in which mothballing and shutdown are combined into no-production. The two-choice model can be estimated with equation 4.15.

Both models recognize sunk costs and allow sunk costs to vary with lease characteristics. The models can be used to test whether there are sunk costs and whether sunk costs are changing in lease characteristics. Allowing for serially-correlated noise and random effects, the models can be estimated with a method of simulated moments estimator.

Chapter 5

DISCUSSION OF THE DATA

This chapter discusses the data and provides some summary statistics of the same. The data set that is used for estimating the empirical model is newly constructed from an oil and gas database for Oklahoma. The construction of the data set and key assumptions underlying the construction are explained in this chapter. Summary statistics of the data are presented, some of which reveal patterns that support the sunk cost hypothesis and the model in Chapter 4.

5.1 The Data Set

The data used in this dissertation has been provided by the Sarkeys Energy Center at the University of Oklahoma. The regression data is based on a subset of the leases and wells in the Sarkeys Energy Center database. The well history and lease production data for a distinct geographic region is organized into an unbalanced, annual panel at lease level. The data set describes the production, mothballing, and shutdown histories of oil leases over the period 1980 to 1997 and includes 18306 observations. The data set indicates the number of wells on a lease, the lease's age, the average well depth, the amount of oil produced, the current production status, and the past production status for each lease in each year.

5.1.1 The Original Data Base

With the cooperation of the Sarkeys Energy Center, I have gained access to a database that contains all oil and gas well, and oil and gas, lease records in Oklahoma that have been transferred to electronic files. This database is comprised of a well-history file and a lease file. The well-history file provides the complete investment history of over 439,000 individual wells, with observations from the early 1900's to the present. The primary source for the well history data are the mandatory reports for well completions, re-completions, workovers, re-entries, and other physical changes to the well that the owner must file with the Oklahoma Corporate Commission.[28] Plugging data, which is based on a well-plugging report, is available for some counties in Oklahoma. Based on the well-history file, it is possible to construct time-series that indicate when a well was drilled and re-drilled, which years it was in operation, and in which year the well was officially plugged. Moreover, this data set provides information on a well's characteristics such as its depth and type.

The lease data file contains information on the oil and gas producing leases. A lease is defined as an oil and gas producing unit, a confined geographic area with one or more wells in place. The primary sources for this data file are reports filed with the Oklahoma Tax Commission. The file includes monthly oil, condensate, associated and non-associated gas production for each lease in Oklahoma for the years 1979 to 1999. In addition, it provides lease characteristic such as the lease name, approximate size, and geographic location.

5.1.2 A Subset of Sarkeys Energy Center's Well and Lease Database

Because of the vast amount of data in the original database, it is necessary to restrict the focus of the empirical investigation to a subset of it. I explore a geographic area that has above average drilling activity in the last 20 years and, thus, appears to be more profitable than most

[28] These reports are typically filed using Form 1002A of the Oklahoma Corporate Commission.

other oil producing regions in Oklahoma.[29] At the same time I observe well closings in this area as elsewhere in Oklahoma.

There are several reasons for using such a area of Oklahoma. In order to identifying the entry and exit thresholds, it is necessary that the latent profits and the latent lease values cross entry and exit thresholds as market conditions change. It is important that changes in market conditions generate changes in the production status of leases. Thus, this area, in which the number of well drilling and closings responds to changes in the price of oil, appears to be suitable for estimating the oil production models developed in Chapter 4.[30]

Additionally, if there is state dependence of mothballing and shutdown in a profitable region, than there should be even more state dependence in an oil producing region of lower profitability where oil producers are less likely to switch back to production after being mothballing or shutdown. Therefore, if I find that the lagged production status of a lease matters in a region with higher drilling activity, it should be even more so in a region with low drilling activity.

According to Sarkeys Energy Center, the most complete well records in Oklahoma are available for the counties: Beaver, Caddo, Carter, Cimarron, Creek, Pontotoc, Stephens, and Texas.[31] Looking at the aggregate number of well drillings for these counties from 1983 to 1998, the number of new well drillings in the counties Beaver, Caddo, Cimarron, Creek, and Pontotoc is relatively low during most of the 90's. This can indicate that most oil fields in these counties are quite depleted. Considering the possibility that shutdown in these counties is driven by the high

[29] Many areas of Oklahoma show little to no drilling activity during the past 20 years.
[30] I failed estimating a similar model as in this dissertation using a data set containing all oil wells in Alaska. Changes in the market conditions in the mid-80's to mid-90's did not generate a sufficient number of well opening or well closings to estimate an empirical model of oil wells drilling and plugging in Alaska.
[31] This is based on the data available in December 1999.

depletion rate rather than low oil prices, I chose to use data from the counties Carter, Stephens, or Texas.

The largest field in the Carter and Stephens counties is Sho-Vel-Tum with 96,480 acres and 6290 oil and gas leases. Additionally, Sho-Vel-Tum is the oil field with the highest production rate and largest proven reserves in Oklahoma in 1999. The field was discovered in 1906.[32] Sho-Vel-Tum appears to be an old, well-explored oil field that is still quite productive. Unfortunately, using all the data available for this field is computationally not feasible. The data used in this dissertation is based on a subset of data for Sho-Vel-Tum, consisting of all well and lease data for the geographic area of Sho-Vel-Tum that is part of Stephens County.[33]

For the above specified geographic region, there are 10,791 well history observations that identify 7,436 wells. The number of leases on record for this area is 3,034 and these are matched with 37,328 records containing the monthly production rates for a lease in a given year.

5.1.3 Construction of the Regression Data Set

The goal is to create a panel of oil leases that includes the current and lagged production state, lease characteristics that change over time, and macro variables. The hardest tasks involved in constructing this panel data set proved to be creating individual well level panel data. The Sarkeys Energy Center's well-history file describes the important dates in a well's life. Depending on a well's history there are one or more observations. Each observation carries important dates such as the drill start, drill end, first production, completion, recompletion, workover, deepening start, deepening end, plug start, and plug end data. Although activities after

---

[32] See Department of Energy Information Administration (2000).
[33] More precisely, this is all the well and lease data for the township and range combinations: 1N5W, 1N4W, 1S5W, 1S4W, 2S5W, and 2S4W.

the initial drilling, for example a workover, generates another 1002A form, it unfortunately does not imply that Sarkeys Energy Center adds a new observation to its database. In many cases, the procedure is to go into the old record(s) and simply write the new dates into an existing observation. For example, if there are several observations for a well, once the well is plugged the plugging date is written into all its records. Or, if the ownership of the well changes, the old owner's name is erased and replace by the new name. This makes it cumbersome to sort the well records and generate a well panel that describes the investment history in a statistically suitable form. With the invaluable help of Kathryn E. Hines at Sarkeys Energy Center, I was able to clean and sort most of the well history data such that it was possible to generate suitable well panel data.

Each well in the U.S. is identified by its unique American Petroleum Institute (API) number. Moreover, Sarkeys Energy Center provides a production unit (PU) number that allows matching wells and leases. The first manipulation in the well history data is to delete all observations for which the API number is missing, and to delete all wells that never carry a PU number. [34] This leaves 7471 well observations that identify 5150 distinct wells on 1710 leases. Of these wells 3455 have only one data entry describing their history. It is easy to organize these observations into a panel. However, the other 1695 wells have up to seven data entries. After cleaning and sorting the multiple well observation data, and deleting wells whose records do not

---

[34] Only wells that can be matched to a lease and vice versa are used for constructing the data set. If there is no well data the lease's investment history is unknown. For wells without matching lease data I do not know if a lease stopped producing before 1979 or if lease data is missing. Not using these observations implies that the data contains only leases that produce at some point during the years 1979 to 1997. Thus, there are no wells in the data set that were plugged prior to 1979 and do not come back to production. This implies that the in-sample probability of re-entry by shutdown leases is higher than it is for the population of all leases. However, if my estimates find that the lagged shutdown status is significant, then, it would be even more significant if the same model were estimated with the true population of leases and their respective production rates.

make sense, there are 5011 wells on 1571 leases in the data set.[35] Based on this data, annual time-series are constructed that indicates when a well existed, when it was in production, its depth, its age, etc.

The lease level production data requires only little manipulation. The original production data is monthly and it is aggregated to the annual level. Moreover, a dummy variable that indicates mothballing is created. Mothballing is defined as: in a given year there is no oil or gas production for at least six consecutive months. This annual production data is then merged with the well panel. Some wells cannot be matched to a lease and vice versa, so that after merging there are 3919 wells on 1361 leases for the years 1979 to 1999. Note that not all wells or leases exist for the entire span of eighteen years. Moreover, this data set is at the well level with lease level production data matched to wells. With respect to the individual well data, it needs to be aggregated to the lease level.

The data is aggregated to the lease level by including the number of wells on a lease in a given year, the number of wells by their respective production category, the average well age in a given year, and the average well depth in a given year. Some wells change their lease over time. The wells for which this happens are pooled by the last lease they belong to. Some leases produce only gas or no well on the lease is classified as an oil well. These leases are dropped from the data. Moreover, there are leases whose data does not make sense. For example, a lease may produce oil although all wells on the lease are plugged according to the well data. These lease are dropped from the data, too.

Based on this data it is possible to define for each year whether a lease is in production, mothballed, or shutdown. A lease is categorized as producing if production took place for five months or longer in a given year; otherwise it is categorized as non-producing. There are two

---

[35] For example, observations with a plugging date earlier than the first drilling date or production

types of non-production states: mothballing and shutdown. A lease is categorized as a shutdown if 1.) all wells on the lease have been plugged, 2.) the lease has been non-producing for two consecutive years, or 3.) the lease is non-producing and stays non-producing throughout all remaining years. The first criterion is the natural definition of shutdown in the context of an oil lease. The second and third criterion are an empirical approximation of the first one. Many leases in the data do not produce for several years, or go out of production and do not return to production through the end of the data period. By law wells that do not produce for a year are required to be plugged. Since the plugging data is not complete, this is the best possible approximation.

The first year of data, 1979, cannot be used in the regression analysis. It is needed to construct the first set of dummy variables that indicate a lease's lagged production state. The same of course is true for the first observation of leases that appear for the first time in the data after 1979. The last two years of data, 1998 and 1999, seem to be very incomplete. The number of well drillings, workovers, deepenings, and pluggings are extremely low when compared to the other years. Most likely the well-history file had not yet been updated completely with the new records. Therefore, the last two years of data are dropped from the regression data set. The final panel of leases has 18306 annual observations for 1095 leases for the years 1980 to 1997.

This lease panel is complemented with a set of market-level variables. These include the average annual real spot price of oil, the annual variance of the spot price, and the annual average of the one-year futures price of crude oil. The average annual real oil price is calculated as the annual average of the monthly spot oil price for West Texas Intermediate deflated by the seasonally-adjusted producer price index for intermediate energy goods.[36]

---

data are dropped.

[36] I use the monthly spot price data for West Texas intermediate [$/bbl of oil] as provided by the Federal Reserve Bank of St. Louis and discount it with the seasonally adjusted, monthly producer

5.2 Descriptive Statistics

Subsection 5.2.1 and 5.2.2 present some descriptive statistics of the data. Section 5.2.1 discusses the lease and oil price data and highlights some of the key features. Section 5.2.2 discusses transition patterns in the data.

5.2.1 Descriptive Statistics of the Data

The final regression data set is an unbalanced panel of oil leases for the years 1980 to 1997. The individual lease characteristics, depth and age, are the average depth and average age of all the wells on the lease in a given year. Depth and age are summarized in *Table 5–1*, Columns 2 and 3 for each year. The average depth is 4577 feet and it stays relatively constant over time. The maximum and minimum observed depth are 365 and 15500 feet. Obviously the average well age increases over time. It is 24.46 years in 1980 and 35.79 years in 1997. However, it does not increase by one for each year because new wells are entering the data set. The oldest average well age in 1980 is 64 years on a lease with seven existing wells. These wells were drilled in 1916.

Column 4 and Column 5 of *Table 5–1* present the average oil production per lease and well. Oil production declines over time. With the exception of 1982, lease and individual well productivity are strictly decreasing. For example, while a lease produces on average 14,019 bbls of oil in 1980 it produces only 5646 bbls of oil in 1997. During the same time period, average well productivity falls from 12.52 bbls/day to 5.29 bbls/day. This pattern is similar to the well-documented decline of oil production and well productivity in the continental U.S.

---

price index for intermediate energy goods (base year=1982), provided by the same data source. The oil futures price is based on Crude Oil, WTI / Global Spot futures, traded at the New York

Columns 6 to 10, *Table 5–1*, present oil price data. The difference between the highest and the lowest real oil price is larger than for the nominal oil price. The simple correlation between the two series is .981, indicating that most movements in the real price result from movements in the nominal price. The relative ordering of prices is not that different. With either price series, the years 1986, 1988, 1994, and 1995 are the years with the lowest price and the years 1980 through 1983 are the years with the highest price.

Column 8 shows the One-Year ahead futures price of oil. In general, the futures price and the nominal oil price are of similar magnitude. The correlation between the nominal price and the futures price is .978. The last column represents the within-year standard deviation of the natural logarithm of the monthly real oil price, a potential measure of oil price uncertainty. Oil price volatility differs from year to year. It is at is highest level in 1990, the year when Iraq invaded Kuwait.

The production states are further disaggregated in *Table 5–2*. Columns 2 through 6 summarize the number of leases by the states: production, production & drilling, production & plugging, mothballing, and shutdown. The states production & drilling, and production & plugging can be interpreted as capacity expansions and capacity contractions. The number of leases with well drillings and the real price of oil are highly correlated with a correlation coefficient of 0.876. The numbers of leases that occupy the states production & plugging, mothballed, or shutdown are negatively correlated with the price of oil.

Wells can be of the type oil, service, gas, or unknown. Columns 7 to 10 of *Table 5–2* display the average numbers of wells on a lease. For example, Column 7 shows the average number of wells on a lease regardless of their type for each year. Generally, the number of wells on a lease does not change very much. The number of service wells does not change over time,

---

Mercantile Exchange as provided by Bridge/CRB.

indicating that secondary recovery methods are not used more extensive as the oil fields depletes more and more. The number of oil wells on a lease is slightly decreasing. The maximum numbers of all wells, production wells, service wells, and gas wells on a lease are 32, 27, 8, and 5, respectively.

Table 5–3 and *Table 5–4* summarize the average depth, age, lease productivity, oil prices, and the number of wells conditional on the lease status in a given year.[37] *Table 5–3* conditions on the states production and no-production. It shows that producing leases have deeper and younger wells. Moreover, there are more wells on producing leases. Producing leases are of higher productivity than non-producing leases.

Table 5–4 distinguishes between production, mothballing, and shutdown. It reveals that leases that are in mothballed state are of lower than average depth. On average, mothballed leases are much less deep than producing or shutdown leases. However, mothballed leases are on average of higher productivity than shutdown leases. *Table 5–3* and *Table 5–4* show that there are differences in key lease characteristics between leases that are in different states in a given year.

Table 5–5 and *Table 5–6* show the average depth, age, lease productivity, and oil price data conditional on lease transitions. The first column in each table displays the state in period *t-1* and the second column displays the state in period *t*. Consider row 2 and 3 in *Table 5–5*, the transition probabilities of a producing lease. Leases that switch from production to no-production are less deep and less productive. Moreover, switches to no-production occur mainly during times of lower oil prices and slightly lower oil price volatility. Next, consider rows 4 and 5, the transition of leases that are in the no-production state. Leases that switch back into production

[37] Leas potential productivity is measured as the last observed strictly positive production rate, i.e. the amount of oil produced in the last year the lease is active.

are less deep, younger, and more productive than leases that stay in the no-production state. The price of oil is higher and oil price volatility lower when leases return to production.

Theoretical models of investment under uncertainty imply that firms are less likely to switch their production status if uncertainty and sunk costs are high. Assuming that the well depth and the standard deviation of the price of oil are measures of sunk costs and uncertainty, columns 3 and 8 provide some evidence that this theoretical implication is right. For further evidence, the last column of *Table 5–5* shows the product of oil price volatility and average depth. This combined measure of sunk costs and uncertainty is on average lower for leases that switch their production status, too.

Table 5–6 differentiates between mothballing and shutdown. Consider leases that produce in *t-1*. Leases that switch to mothballing are less deep and less productive than leases that stay in production. Surprisingly, leases that shut down are deeper and slightly more productive than leases that switch to mothballing. However, the price of oil is lowest and the average age is highest for leases that switch from production to shutdown. When considering leases that are mothballed in *t-1*, the only clear picture that emerges is that age, the standard deviation of the oil price, and standard deviation of the oil price combined with depth are increasing through the states production, mothballing, and shutdown in period *t*. Finally, leases that return to production after being shut down in *t-1* are younger and more productive than the ones that stay shut down. On average, leases return to production during times of higher oil prices and lower uncertainty.

Table 5–5 and *Table 5–6* reveal some expected relations. Nevertheless, some patterns are not expected. For example, conditional on being mothballed in *t-1*, leases that return to production are less productive than leases that stay mothballed or switch to shutdown. However, this analysis cannot reveal how different combinations of the explanatory variables affect the production, mothballing, and shutdown decision. The empirical analysis in Chapter 6 will provide more light on this matter.

5.2.2 Transition Patterns Between Production, Mothballing, and Shutdown

The theoretical model of production, mothballing, and shutdown predict that the combination of sunk costs and uncertainty create state dependence in the production and investment history. For example, an oil lease that currently produces has already sunk some costs for being in the producing state and, thus, is more likely to produce next period. Similarly, an oil lease that is shut down has to sink costs for returning to production. Thus a shutdown lease is expected to stay in the shutdown state with high probability. The transition rates displayed in *Table 5–7* and *Table 5–8* provide evidence that this conjecture is correct. The tables describe, for each year, the probability that a lease moves from a particular production state in *t-1* to a production state in period *t*.

Table 5–7 summarizes the transition probability for the oil lease data set when only the production and no-production are possible states. The third and forth columns represent the probabilities of staying in production and switching from production to no-production, respectively. Column five displays the probability of switching from no-production to production, and column six contains the probability of being in the no-production state for another period.

Providing evidence for the importance of sunk costs, the probabilities of staying in the lagged production state are very high. In any year at least 90% of the leases in operation continue production in the following year. For example, 92 percent of the leases that are producers in 1993 stay in production in 1994, the year with the lowest real and second lowest nominal price of oil. Overall, the production state is very persistent. Although the real oil price fluctuates between

$41.65 and $14.77 in the time period 1980 to 1997, the transition rate from production to production is always between 99.1 and 91.2 percent. It appears that very large price changes are necessary to induce a lease to switch to no-production. This is even more true when taking into account that the years of lower state persistence are towards the end of the data period when oil production rates are relatively low, see *Table 5–1*.

The state no-production is less persistent. The transition rates for staying in the state no-production are between 64.7 and 92.7 percent. Note that the persistence of no-production is increasing with time. This phenomenon probably reflects decreasing lease profitability because of lower lease productivity and lower oil prices, which coincide towards the end of the data period.

The no-production state can be divided into the categories mothballed and shutdown. *Table 5–8* summarizes the annual transition probabilities when the production, mothballing and shutdown states are observed. The sunk costs involved with switching between the categories production, mothballed and shutdown depend on the lagged state. For example switching from production to mothballing is less costly than switching from production to shutdown. This should be reflected in the transition rates: state persistence should be increasing in sunk costs.[38]

Columns 2, 3, and 4 of *Table 5–8* present the transition rates of switching from production in period *t-1* to the production, mothballing, and shutdown state in period *t,* respectively. Columns 5 through 9 are similar statistics for the states mothballing and shutdown in *t-1*. The transition rates from production to production indicate state dependence and are, of course, the same as in *Table 5–7*. Until 1993 leases that stop production are more likely to mothball than to shutdown. During this time the conditional number of switches to shutdown is

---

[38] See equation 4.7. The weight attached to the dummy variable for lagged shutdown is larger than the one attached to the dummy variable for lagged mothballing.

relatively low and in most years below 1 percent. On average, only 2.3 and 2 percent of producing leases switch to mothballing or shutdown, respectively, in any year.

The picture for transitions conditional on mothballing is somewhat different. A large percentage of mothballed leases switch to either production or shutdown. In fact, the average transition rates from mothballing to production, mothballing, or shutdown are 40.5, 37.8, and 21.7 percent, respectively. Nevertheless, the large switching probabilities make sense for the mothballing state as it is the intermediate state for temporary non-production that can be utilized for a limited amount of time only.

Leases in the shutdown state switch either to production or they stay shutdown. They never switch to mothballing. On average, 80 percent of the leases that were shutdown stay shutdown the next year. Shutdown is much more persistent than mothballing. In general, persistence in the shutdown state is increasing over time. The lowest persistence is in 1984 with 61 percent, the highest is recorded in the last year of data with 95.8 percent.

Recall that categorizing non-producing leases into mothballing and shutdown is somewhat arbitrary; for explanation see Section 5.1.3. In this data set, all leases that are out of production for two consecutive years are classified as shutdown. Unfortunately, some of these leases might just be mothballed and their well(s) are not plugged. Since mothballed leases are more likely to return to production than shutdown leases this might increase the transition rate from shutdown to production and lower the state persistence of shutdown. In any case, if there is indeed such a bias in the data, it lowers the importance of the lagged shutdown state and, thus, any sunk cost estimate. Findings regarding the significance of sunk costs based on the lagged shutdown state still hold for the “true” data.

5.3 Summary

The transition patterns in the data reveal that the states production and shutdown are very persistent. The transition patterns provide first evidence that sunk costs play an important role in the decision to produce oil. Moreover, there is some evidence that the decision to produce, mothball, or shutdown is affected by variables such as well depth, well age, lease productivity, the price of oil, and oil price volatility.

Table 5–1: Average Depth, Age, Production Rates, and Prices by Year

Year	Well Depth [feet]	Well Age	Oil Production [bbls/year], lease	Oil Production [bbls/day], well	Oil Price [$/bbl]	Real Oil Price [real $/bbl]	One-Year Future	Ln(Real Oil Price)	S.D. of ln(Real Oil Price)	Number of Leases
1980	4577	24.46	14019	12.52	37.38	41.65	32.54	3.73	0.064	884
1981	4560	24.74	12368	10.81	36.67	37.42	37.67	3.62	0.044	911
1982	4548	25.22	12519	10.84	33.64	33.63	32.11	3.51	0.066	931
1983	4533	25.62	10412	8.95	30.40	30.02	28.83	3.40	0.034	954
1984	4512	26.19	9808	8.34	29.28	28.24	28.90	3.34	0.052	970
1985	4566	26.61	9050	7.76	27.97	27.12	25.61	3.30	0.048	992
1986	4570	27.40	8729	7.50	15.04	14.99	14.90	2.69	0.162	1000
1987	4577	28.02	7905	6.81	19.16	18.64	18.30	2.92	0.056	1014
1988	4564	28.75	7832	6.76	15.96	14.94	15.93	2.70	0.096	1024
1989	4545	29.30	7516	6.50	19.59	17.45	17.53	2.86	0.047	1040
1990	4548	29.84	7194	6.30	24.49	20.97	21.93	3.02	0.243	1058
1991	4542	30.70	7417	6.54	21.48	18.43	20.12	2.91	0.065	1064
1992	4556	31.55	6936	6.20	20.56	17.54	20.12	2.86	0.057	1070
1993	4575	32.46	6534	5.94	18.46	15.52	19.48	2.74	0.097	1073
1994	4585	33.40	6175	5.64	17.19	14.26	17.39	2.65	0.094	1075
1995	4629	34.18	5990	5.56	18.43	14.77	17.54	2.69	0.044	1082
1996	4691	34.84	5753	5.38	22.15	17.34	18.52	2.85	0.093	1093
1997	4702	35.79	5646	5.29	20.60	16.14	19.97	2.78	0.076	1095
Mean	4577	29.39	8433	7.42	23.80	22.17	22.63	3.03	0.080	1018

The annual averages of oil production are based on the lagged (*t-1*) oil production of all leases that are not mothballed or shutdown.
Oil production per well is based on lease oil production and the number of wells that are classified as producing oil wells.
One-year futures were not traded during the years 1980, 1981, and 1982. The implied One-Year Futures Price is approximated from 9-month futures.

Table 5–2: Lease State and Well Types by Year

Lease State and Well Types by Year									
	Number of Leases in State:					Average Number of Wells per Lease, by Well Type:			
Year	Production	Production & Drilling	Production & Plugging	Mothballed	Shutdown	All Types	Oil	Service	Gas
1980	708	56	6	14	100	3.07	2.48	0.340	0.080
1981	724	74	2	17	94	3.14	2.57	0.340	0.083
1982	776	51	3	13	88	3.16	2.59	0.343	0.087
1983	766	65	6	28	89	3.19	2.60	0.352	0.090
1984	796	56	7	39	72	3.22	2.63	0.366	0.096
1985	832	41	5	27	87	3.19	2.61	0.366	0.096
1986	852	13	2	37	96	3.19	2.61	0.372	0.095
1987	829	33	8	42	102	3.18	2.60	0.368	0.097
1988	863	24	4	39	94	3.17	2.60	0.366	0.099
1989	840	38	4	44	113	3.17	2.58	0.374	0.102
1990	858	39	8	30	122	3.13	2.53	0.382	0.105
1991	840	16	14	36	157	3.11	2.52	0.378	0.103
1992	809	24	27	42	166	3.07	2.50	0.365	0.101
1993	822	6	19	30	194	3.01	2.46	0.356	0.100
1994	799	7	6	42	219	3.00	2.45	0.358	0.098
1995	768	7	18	33	253	2.95	2.41	0.344	0.097
1996	739	13	8	18	310	2.93	2.40	0.341	0.097
1997	734	5	6	18	326	2.92	2.39	0.340	0.096
Mean	798	32	9	31	149	3.10	2.53	0.358	0.096

Table 5–3: Averages Conditional on the States: Production and No-Production

Averages Conditional on the States: Production and No-Production											
Lease State in Year *t*	Depth [feet]	Well Age	Oil Production [bbls/year], lease	Oil Price	Real Oil Price	S.D. of ln(Real Oil Price)	S.D. of ln(Real Oil Price)* Depth	Number of Wells per Lease, by Well Type:			
								All Types	Oil	Service	Gas
Production	4639	29.00	9382	23.70	22.07	0.083	0.386	3.199	2.618	0.364	0.102
No-Production	4292	32.40	3068	22.31	19.96	0.084	0.354	2.619	2.100	0.334	0.067

The oil production data for leases in the no-production state is based on the last production rate when being in the production state.

Table 5–4: Averages Conditional on the States: Production, Mothballing, and Shutdown

Averages Conditional on the States: Production, Mothballing, and Shutdown

Lease State in Year *t*	Depth [feet]	Well Age	Oil Production [bbls/year], lease	Oil Price	Real Oil Price	S.D. of ln(Real Oil Price)	S.D. of ln(Real Oil Price)* Depth	Number of Wells per Lease, by Well Type:			
								All Types	Oil	Service	Gas
Production	4639	29.00	9383	23.70	22.07	0.083	0.386	3.199	2.618	0.364	0.102
Mothballed	3349	28.86	3550*	22.23	20.40	0.083	0.272	2.909	2.297	0.326	0.095
Shutdown	4485	33.12	2969*	22.32	19.87	0.084	0.370	2.560	2.060	0.336	0.062

* The oil production data for mothballed and shutdown leases is based on their last production rate when being in the production state.

Table 5–5: Averages Conditional on State *t* and State *t-1* for Production and No-Production

Averages Conditional on State *t* and State *t-1* for Production and No-Production									
State in *t-1*:	State in *t*:	Depth [feet]	Well Age	Oil Production [bbls/year], lease	Oil Price	Real Oil Price	S.D. of ln(Real Oil Price)	S.D. of ln(Real Price) *Depth
Production	Production	4656	29.42	9537	23.67	22.03	0.083	0.388
	No Production	3974	29.88	3427*	21.45	19.13	0.082	0.319
No Production	Production	4278	19.92	6068	24.43	23.04	0.080	0.339
	No Production	4374	33.04	2977*	22.52	20.18	0.084	0.363

* The oil production data for the leases with zero production is based on their last production rate when being in the production state.

Table 5–6: Averages Conditional on State *t* and State *t-1* for Production, Mothballing, and Shutdown

Averages Conditional on State *t* and State *t-1* for Production, Mothballing, and Shutdown

State in *t-1:*	State in *t:*	Depth [feet]	Well Age	Oil Production [bbls/year], lease	Oil Price	Real Oil Price	S.D. of ln(Real Oil Price)	S.D. of ln(Real Price) *Depth
Production	Production	4656	29.42	9537	23.67	22.03	0.083	0.388
	Mothballed	3273	28.42	3039*	22.13	20.42	0.081	0.252
	Shutdown	4784	31.57	3875*	20.66	17.63	0.083	0.396
Mothballed	Production	3159	27.21	3022	21.44	19.13	0.085	0.264
	Mothballed	3486	29.65	4472*	22.39	20.37	0.088	0.309
	Shutdown	3414	32.69	3413*	21.59	19.37	0.099	0.358
Shutdown	Production	4797	16.54	7482	25.81	24.85	0.078	0.374
	Shutdown	4502	33.35	2823*	22.59	20.20	0.083	0.367

* The oil production data for the leases with zero production is based on their last production rate when being in the production state.

Table 5–7: Transition Rates for Production and No-Production by Year

Transition Rates for Production and No-Production by Year

Year *t*	State in *t-1*: / State in *t*:	Production (**P**)		No-Production (**NP**)		Number of Leases in Year *t-1* that are:		Real Oil Price
		P	**NP**	**P**	**NP**	**P**	**NP**	
1980		0.980	0.020	0.331	0.669	736	148	41.65
1981		0.991	0.009	0.262	0.738	770	141	37.42
1982		0.985	0.015	0.321	0.679	800	131	33.63
1983		0.957	0.043	0.347	0.653	830	124	30.02
1984		0.970	0.030	0.353	0.647	837	133	28.24
1985		0.977	0.023	0.293	0.707	859	133	27.12
1986		0.960	0.040	0.197	0.803	878	122	14.99
1987		0.962	0.038	0.245	0.755	867	147	18.64
1988		0.956	0.044	0.383	0.617	870	154	14.94
1989		0.943	0.057	0.284	0.716	891	148	17.45
1990		0.981	0.019	0.229	0.771	882	175	20.97
1991		0.935	0.065	0.152	0.848	905	158	18.43
1992		0.952	0.048	0.166	0.834	869	199	17.54
1993		0.943	0.057	0.171	0.829	860	211	15.52
1994		0.920	0.080	0.146	0.854	847	226	14.26
1995		0.946	0.054	0.094	0.906	812	267	14.77
1996		0.912	0.088	0.131	0.869	791	297	17.34
1997		0.950	0.050	0.073	0.927	759	330	16.14
Mean		0.957	0.043	0.232	0.768	837	180	22.17

Table 5–8: Transition Rates for Production, Mothballing, and Shutdown by Year

Transition Rates for Production, Mothballing, and Shutdown by Year													
Year *t*	State *t-1:*	Production (**P**)			Mothballing (**M**)			Shutdown (**S**)		Number of Leases in Year *t-1* that are:			Real Oil Price
	State *t:*	**P**	**M**	**S**	**P**	**M**	**S**	**P**	**S**	**P**	**M**	**S**	
1980		0.980	0.019	0.001	-	-	-	0.331	0.669	736	0	148	41.65
1981		0.991	0.009	0.000	0.333	0.667	0.000	0.254	0.746	770	15	126	37.42
1982		0.985	0.011	0.004	0.471	0.235	0.294	0.298	0.702	800	17	114	33.63
1983		0.957	0.028	0.016	0.385	0.385	0.231	0.342	0.658	830	13	111	30.02
1984		0.970	0.023	0.007	0.214	0.714	0.071	0.390	0.610	837	28	105	28.24
1985		0.977	0.017	0.006	0.308	0.308	0.385	0.287	0.713	859	39	94	27.12
1986		0.960	0.031	0.009	0.259	0.370	0.370	0.179	0.821	878	27	95	14.99
1987		0.962	0.032	0.006	0.432	0.378	0.189	0.182	0.818	867	37	110	18.64
1988		0.956	0.028	0.016	0.476	0.357	0.167	0.348	0.652	870	42	112	14.94
1989		0.943	0.035	0.022	0.474	0.342	0.184	0.218	0.782	891	38	110	17.45
1990		0.981	0.014	0.006	0.295	0.409	0.295	0.206	0.794	882	44	131	20.97
1991		0.935	0.034	0.031	0.533	0.167	0.300	0.063	0.938	905	30	128	18.43
1992		0.952	0.026	0.022	0.417	0.528	0.056	0.110	0.890	869	36	163	17.54
1993		0.943	0.026	0.031	0.429	0.190	0.381	0.107	0.893	860	42	169	15.52
1994		0.920	0.033	0.047	0.400	0.467	0.133	0.107	0.893	847	30	196	14.26
1995		0.946	0.021	0.033	0.366	0.390	0.244	0.044	0.956	812	41	226	14.77
1996		0.912	0.013	0.076	0.485	0.242	0.273	0.087	0.913	791	33	264	17.34
1997		0.950	0.017	0.033	0.611	0.278	0.111	0.042	0.958	759	18	312	16.14
Mean		0.957	0.023	0.020	0.405	0.378	0.217	0.200	0.800	837	29	151	22.17

148 leases are non-producing in 1979 and it is not possible to tell whether they are mothballed or shutdown. The vast majority of these leases does not take up production for several years and these leases are classified as shutdown. Because of this, transition rates conditional on "mothballing" in 1980 are not available.

Chapter 6

EMPIRICAL RESULTS

This chapter estimates the oil production models that are developed in Chapter 4. Results are reported for a model with two states, production and no-production, and a three-state model that distinguishes between production, mothballing, and shutdown.

6.1 Data and Variables

The models in Chapter 4 are estimated using the data described in Chapter 5. Some of the variables have been slightly manipulated for the regression analysis. The variable Age is the average age of the wells on a lease divided by ten. The variable Depth is the average well depth on a lease in thousands of feet. Oil production is the last, lagged year of strictly positive oil production data in thousands of barrels.[39]

Recall that the model is estimated using macro variables $\boldsymbol{P}_t$, lease characteristics $\boldsymbol{X}_{it}$, and variables that proxy entry and exit costs $\boldsymbol{H}_{it}$. The macro variables $\boldsymbol{P}_t$ are the logarithm of the real price of oil and the within-year standard deviation of the prior. The vector $\boldsymbol{X}_{it}$ contains oil production data, the average well depth on a lease, the average well age on a lease, the number of wells on a lease, and the number of service wells on a lease. To correct the initial conditions problem, an indicator of the lease status in period in the first year of observations is also

[39] The last amount of oil that was produced is a good indicator for the production potential of a lease.

included.[40] The variables that affect the entry and exit costs, H_{it}, are a constant, the average well depth (Depth), and the number of wells on a lease (# of Wells).

The variable S.D. Price, the within-year standard deviation of the oil price, is assumed to proxy the level of uncertainty. A year with high oil price volatility is a year in which future oil prices are more uncertain than in years with low oil price volatility. This variable is expected to affect the decision to change the production status through the option value. The average well depth on a lease is used as a measure of sunk costs, since drilling and workover costs are increasing in depth. Moreover, sunk costs are expected to be increasing in the number of wells on a lease.

The results in this chapter are the final regression specification. Preliminary estimations included variables such as the one-year oil futures price, the lease size, and gas production data. However, these variables were not helpful in explaining the observed production decisions. Moreover, since the employed estimation procedure is very costly with respect to computer time, especially for the three-choice model, there is a bias in my regression specification to a parsimonious model.

6.2 Estimates for the Two-Choice Model: Production and No-Production

Table 6–1 reports the results for the binomial model that distinguishes between the choices production and no-production. Recall equation 4.14

---

[40] This variable is included to condition on the lease's initial observed state in the data set. See Wooldridge (2000).

$$np_{it} = \begin{cases} 1 \text{ if } \quad NoP_{it}^{*} - M_i + np_{it-1}(M_i + RM_i) + \varepsilon_{it} \geq 0 \\ \\ 0 \text{ otherwise,} \end{cases} \qquad (\quad 6.1 \quad)$$

where

$$NoP_{it}^{*} \equiv \beta E_t \left[V_{i,t+1}(\Omega_{i,t+1}) | np_{it} = 0 \right] - \beta E_t \left[V_{i,t+1}(\Omega_{i,t+1}) | np_{it} = 1 \right] - \Pi_{it}(\bullet). \qquad (\quad 6.2 \quad)$$

A lease is in the no-production state if the current profit and the future lease value in the no-production state are larger than in the production state. Since switching between states is costly, the current lease value depends on the lagged production state, too.

Controlling for the latent profit and value differences using the macro variables $\boldsymbol{P_t}$ and lease characteristics $\boldsymbol{X_{it}}$ one can test for sunk costs with the empirical estimation equation 4.15:

$$np_{it} = \begin{cases} 1 \text{ if } \alpha + \boldsymbol{\beta P_t} + \boldsymbol{\delta X_{it}} + \gamma np_{it-1} \boldsymbol{H}_{it} + \varepsilon_{it} \geq 0 \\ \\ 0 \text{ otherwises.} \end{cases} \qquad (\quad 6.3 \quad)$$

The relation between the two choices is normalized by the choice production. A negative coefficient estimate implies that the probability of choosing no-production is decreasing in the corresponding variable and the probability of choosing production is increasing in the corresponding variable. Thus, a positive coefficient in $\boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ implies that the latent value of the no-production state is increasing compared to the latent value of the production state. For example, a higher oil price should increase the value of the production state and, thus, its coefficient should be negative. With respect to testing for sunk costs, the coefficients of interest are $\boldsymbol{\gamma}$. The coefficients in $\boldsymbol{\gamma}$ are expected to be positive.

Finally, for estimating the model, it is necessary to specify the error term ε_{it}. A potential source of persistence of the no-production state is serial correlation in the error terms. Thus, to properly estimate the importance of sunk costs, it is necessary to control for serial correlation. This is accomplished using the probit estimator and error structure outlined in section 4.2.1 that allows for serial correlation and a random lease effect. Recall that for the model with the two choices the error structure is

$$\varepsilon_{it} = \sqrt{1-\lambda} \times \omega_{it} + \sqrt{\lambda} \times \upsilon_i \text{ , where } \omega_{it} = \rho\omega_{it-1} + \eta_{it}\text{; } \quad \eta_{it}, \upsilon_i \sim N(0,1). \qquad (6.4)$$

The parameter ρ governs the AR(1) process and the parameter $\lambda \in [0,1]$ measures the fraction of variance that is due to the permanent lease component. *Table 6–1* reports the coefficient estimates and t-statistics for the empirical model in equation 6.3 using the both, the simple probit estimator, where $\lambda = \rho = 0$, and the error specification in equation 6.4. The latter specification is estimated with a method of simulated moment estimator (MSM), described in Section 4.2.2. The choice probabilities for the MSM estimator are simulated with the Geweke-Hajivassiliou-Keane (GHK) simulation algorithm.

Table 6–1, Column 2 to Column 3 report the simple probit results, and column 4 and column 5 report the AR(1)-random effects probit coefficient estimates and t-statistics. Consider the estimates of the probit covariance parameters λ and ρ in lowest two rows. The MSM estimates imply that about 13.6 percent of the variance is due to the permanent lease-level errors. Moreover, the MSM results show that the contemporaneous errors are negatively serially-correlated with ρ being -.244. While the parameters from the more general error structure are statistically significant, they appear to have little effect on the remaining coefficients.

Next, consider the macro variables that control for differences in profit levels, Oil Price and S.D. Price. The probability of choosing no-production is decreasing in the price of oil. As oil production is more profitable when oil prices are high, this is the anticipated result. Oil price volatility is negative and significant, which indicates that the value of being in production is increasing in uncertainty. Because of the option value, even risk neutral firms should be affected by uncertainty. The significant coefficient on S.D. Price, -.980 with the MSM probit estimator, shows that the latent lease values are indeed affected by uncertainty.

The characteristics of the individual lease are also statistically important. A higher past production rate and a larger number of wells on a lease, most likely indicators of lease profitability, increase the probability of production. The probability of no-production is increasing in depth, age, and the number of service wells. Oil productivity declines with cumulative production, e.g. age. In addition, secondary recovery methods that require service wells are expensive. The variable that corrects for initial conditions, I(Initially No Production), is positive and, when estimated with MSM, significant.

The last four rows of the table contain the coefficients that test for sunk costs. The coefficients on the variables that are interacted with the indicator for lagged no-production *np(t-1)* are positive and significant across the two different estimation methods. The estimates for the lagged production indicator are similar with 1.884 and 1.916 for probit and MSM probit, respectively. The lagged production history helps explaining the current production state of a lease, implying that there are sunk costs. Moreover, state dependence is increasing in well depth and the lagged number of wells on a lease. Assuming that sunk costs are increasing in well depth and the number of wells, this is the anticipated result. Depending on the structure of the regression errors, the lagged number of wells on a lease interacted with *np(t-1)* is positive or negative.

In summary, a lease's current production status is heavily affected by its status in the previous period. These results imply that sunk costs are a determinant of the current production status and that the state persistence is increasing with variables like well depth, that proxy sunk costs. Moreover, oil price volatility affects the production status through the approximated latent current profits and future lease value.

6.3 Estimates for the Three-Choice Model: Production, Mothballing, and Shutdown

Table 6–2 reports the estimation results for the model with the choices: production, mothballing, and shutdown. The table estimate the choice model of equation 4.7:

$$\text{produce}: Z_{it}^{1} = 0$$

$$\begin{aligned}\text{mothball}: Z_{it}^{2} &= -vm_i - M_i - \Pi_{it} + \beta\Delta V_{i,t+1}^{e}(m_{it}=1) \\ &+ m_{it-1}(M_i + RM_i) + \varepsilon_{i2t}\end{aligned} \qquad (6.5)$$

$$\begin{aligned}\text{shutdown}: Z_{it}^{3} &= -(S_i + M_i) - \Pi_{it} + \beta\Delta V_{i,t+1}^{e}(s_{it}=1) \\ &+ m_{it-1}(M_i + RM_i) + s_{it-1}(S_i + M_i + RM_i + RS_i) + \varepsilon_{i3t}.\end{aligned}$$

Each production state generates current profits and a discounted future expected lease value. Choice j is observed in period t if $Z_{it}^{j} > Z_{it}^{k}\ \forall\, k \neq j$. For empirical identification, the choice values are normalized by the production state. The current profits from production are Π_{it}. The term $\Delta V_{i,t+1}^{e}(m_{it}=1)$ denotes the differences in the future expected lease value of a mothballed lease and producing lease, and $\Delta V_{i,t+1}^{e}(s_{it}=1)$ denotes the difference in the future expected lease value of a shut down lease and a producing lease. The term M_i+RM_i captures the sunk costs of switching from production to mothballing and back to production. The sunken entry and exit costs for switching from production to shutdown and back to production are S_i+M_i+RM_i+RS_i. A lease chooses the state that gives the highest current profits and future expected lease value. The choice that generates the highest profit and value depends on lease characteristics, market conditions and, because of the sunk switching costs, on the lagged lease status.

The empirical estimation equation of this model is specified in equation 4.9

$$\text{produce}: Z_{it}^{1} = 0$$

$$\text{mothball}: Z_{it}^{2} = \alpha_2 + \boldsymbol{\beta}_2 \boldsymbol{P}_t + \boldsymbol{\delta}_2 \boldsymbol{X_{it}} + \gamma_2 m_{it-1} \boldsymbol{H_{it}} + \varepsilon_{i2t} \qquad (6.6)$$

$$\text{shutdown}: Z_{it}^{3} = \alpha_3 + \boldsymbol{\beta}_3 \boldsymbol{P_t} + \boldsymbol{\delta}_3 \boldsymbol{X_{it}} + \gamma_3 m_{it-1} \boldsymbol{H_{it}} + \boldsymbol{\tau}_3 s_{it-1} \boldsymbol{H_{it}} + \varepsilon_{i3t}.$$

The micro and macro variables $\boldsymbol{P_t}$ and $\boldsymbol{X_{it}}$ are assumed to sufficiently control for the latent profit and value differences. Equation 6.6 is used to test for sunk costs. As pointed out, there should be state persistence in the production status if sunk costs, namely M_i+RM_i and S_i+M_i+RM_i+RS_i, matter for choosing the production status. State persistence, and how state persistence varies with different levels of sunk costs, is captured in the terms $\boldsymbol{\gamma}_2 m_{it-1}\boldsymbol{H}_{it}$, $\boldsymbol{\gamma}_3 m_{it-1}\boldsymbol{H}_{it}$, and $\boldsymbol{\tau}_3 s_{it-1}\boldsymbol{H}_{it}$. For example, the vector $\boldsymbol{H}_{it}$ contains an intercept term and the variable Depth. If sunk costs matter and are increasing in Depth, then the coefficients on these two variables will positive.

The model is estimated with two specifications of error terms. First, I specify logistic errors, which result in the multinomial logit (MNL) estimator. Second, I specify normal errors, which result in a multinomial probit (MNP) estimator. The MNP model is estimated with a method of simulated moments (MSL) estimator. The MNL, a standard estimator for estimating multinomial choice models, serves as a benchmark for the simulation estimator. Although MNL software is readily available, the estimator is less desirable than MNP because of its IIA property.

The normal errors in the MNP model follow the specification in equation 4.16 allowing for serial correlation and random effects:

$$\begin{pmatrix} \varepsilon_{i1t} \\ \varepsilon_{i2t} \end{pmatrix} = \sqrt{1-\lambda} \times \begin{pmatrix} \omega_{i1t} \\ \omega_{i2t} \end{pmatrix} + \sqrt{\lambda} \times \begin{bmatrix} 1 & 0 \\ \psi_{12} & \psi_{22} \end{bmatrix} \begin{pmatrix} \nu_{i1} \\ \nu_{i2} \end{pmatrix}, \text{ where}$$

$$\begin{pmatrix} \omega_{i1t} \\ \omega_{i2t} \end{pmatrix} = \begin{bmatrix} \rho_1 & 0 \\ 0 & \rho_2 \end{bmatrix} \begin{pmatrix} \omega_{i1t-1} \\ \omega_{i2t-1} \end{pmatrix} + \begin{bmatrix} 1 & 0 \\ \phi_{12} & \phi_{22} \end{bmatrix} \begin{pmatrix} \eta_{i1t} \\ \eta_{i2t} \end{pmatrix}, \text{ and} \qquad (6.7)$$

$$\boldsymbol{\eta}_{it}, \boldsymbol{\upsilon}_i \sim N(0,1).$$

As discussed earlier, the parameters ρ_1 and ρ_2 govern the AR(1) process. The parameter $\lambda \in [0,1]$ measures the fraction of variance that is due to the permanent lease component. Correlation across choices is governed by parameters ϕ_{12} and ψ_{12}.

6.3.1 MNL Results

Table 6–2 displays the coefficient estimates and t-statistics for the three-choice model. This section discusses the logit results briefly. Columns 2 to 5 report results for the logistic model. Columns 2 and 3 report the coefficients and t-statistics for choosing the mothballing state. Similarly, columns 4 and 5 report the coefficients and t-statistics for choosing the shutdown state. A positive coefficient estimate implies that the probability of being in the state the coefficient is associated with, either mothballing or shutdown, is increasing in a particular variable. This section discusses the logit results briefly.

First, consider the macro and micro variables that approximate latent profits and future expected lease values. The logit results imply that the probabilities of mothballing and shutdown are decreasing in the price of oil. The higher the price of oil is the higher the probability that a lease is in the production state. S.D. Price is negative and significant for mothballing and shutdown at the 10 percent level. Uncertainty seems to play a role for future expected lease values. The variable Oil Production approximates a lease's productivity and the probabilities of mothballing and shutdown are both decreasing in it. Depth and Depth^2 are significant for the mothballing state only. The probability of mothballing is deceasing in Depth at a deceasing rate. In contrast, the variables Age and Age^2 are only significant for the shutdown state. The probability of shutdown is increasing in Age at a decreasing rate. The probabilities of the two non producing states are decreasing in the number of wells on a lease. However, the number of service wells is insignificant for mothballing and shutdown. The two variables that are included for conditioning the estimation on a lease's initial state are mostly insignificant. Only I(Initially Shut) is significant for shutdown.

The lagged state indicators *m(t-1)* and *s(t-1)* and their interactions with Depth, and the number of wells on a lease will only be significant if the costs involved in mothballing and shutdown are non-recoverable. The mothballing coefficient estimates on *m(t-1)* is 1.591 The shutdown coefficient estimates on *m(t-1)* is 2.332. Thus, there are sunk mothballing costs. The interaction term Depth**m(t-1)* is positive and significant for staying mothballed, but insignificant for switching from mothballing to shutdown. In addition, the probability of staying mothballed and the probability of switching from mothballing to shutdown are both increasing in the number of wells on a lease.

The indicator for lagged shutdown *s(t-1)* and interaction terms are only estimated for the shutdown choice, as there are no switches from shutdown to mothballing. Lagged shutdown is positive and significant with 5.838. The coefficient estimate for *s(t-1)* is larger than the estimates for *m(t-1)*, suggesting that the shutdown sunk costs are larger than the mothballing sunk costs. The shutdown costs appear not to be affected by Depth. The number of wells interacted with lagged shutdown is significant at the 10 percent level.

Section 6.3.2 discusses the probit estimates. The logit and probit results are very similar in the signs and significances of the regression coefficients, and the estimation results are discussed in more depth below.

6.3.2 MNP Results

Columns 6 through 9 in *Table 6–2* d report the estimates and t-statistics for normal errors. The final rows report the estimates for the parameters that govern the structure of the normal errors that is outlined in equation 6.7. For estimating the probit model, the error structure in equation 6.7 had to be restricted. Keane (1992) suggest that this due to an empirical identification problems. The parameters ϕ_{22} and ψ_2 are fixed at $\phi_{22}=\psi_{22}=1$. Moreover, the reported error structure does not allow for auto-correlation and correlation across choices with the parameter specification $\rho_1=\rho_2=\phi_{12}=\psi_{12}=0$. Only the parameter λ is estimated. The estimate of λ is .114 implying that about 14 percent of variance of the regression errors is due to the unobserved lease level random effect.

Attempts to estimate a less restrictive error structure were unsuccessful and rendered unreasonable parameters. Keane (1992) provides an explanation for this estimation problem. He argues that it is hard to disentangle covariance parameters from regressor coefficients in the MNP model. Keane finds that parameter identification in the MNP model without exclusion restrictions, i.e. restrictions that some exogenous variables in the model do not affect the utility level of certain alternatives, is extremely fragile. Keane's Monte Carlo study of a trinomial probit-model without exclusion restrictions results in a close-to-singular Hessian and questionable parameter estimates. Keane concludes that it is important to realize the practical limitations of the MNP estimator. Moreover, he notes that the necessary choice specific variables are not available for many economic applications.

Table 6–2 provides MNP estimates for the model where the covariances across choices and autocorrelations across time are restricted to zero. The first 12 rows of coefficient estimates report the results for the intercept, the macro variables $\boldsymbol{P_t}$, and the micro variables $\boldsymbol{X_{it}}$. The next 6 rows provide the results regarding sunk costs and display the estimates for m_{it-1} and s_{it-1} and their interactions with other variables. Note first that the signs and significance of the estimated coefficients are relatively robust to the logit and probit specification. For simplicity, unless otherwise indicated, the reminder of this discussion refers to the MSM Probit results.

Consider the variables that control the latent profit and future expected lease value difference between mothballing and production, and shutdown and production. The probabilities of mothballing and shutdown are decreasing in the price of oil. This implies that the profit and the lease value differences of mothballing and production, and shutdown and production are increasing in the price of oil. The value of the production state increases more than the values of the mothballing and shutdown states when the price of oil rises. Oil price volatility, S.D. Price, has a significant negative coefficient for mothballing and shutdown. I conclude that the current level of uncertainty affects the future expected lease value. Theoretical real options models suggest that uncertainty affects the value of the real investment option. The estimate provides evidence for this hypothesis.

How do lease characteristics affect the probabilities of being in the mothballed or shutdown state? The variable Oil Production, approximating a lease's production potential, is negative and significant for the shutdown state. Oil Production is negative and significant at the 10 percent level for the mothballing state. Thus, the higher a lease's production potential is the lower its probability of being shutdown or mothballed. The variable Depth has a negative impact and the variable Depth^2 has a positive impact on the probability of being mothballing. Accordingly, deeper wells are less likely to be mothballed at a decreasing rate. Depth and Depth^2 are insignificant for the shutdown state.

The variable Age is insignificant for the mothballing state. Age, however, has a positive impact on the probability of shutdown. Age^2 has a negative impact on the shutdown state. Combining this, the probability of observing the shutdown state is increasing at a decreasing rate in the average well age on a lease. Oil reserves are decreasing in cumulative production or age. Age can be interpreted as an indicator of the future production potential. Thus, it makes sense that shutdown is positively affected by Age, and that it is not that important for mothballing.

The larger the number of wells on a lease, the lower the probability of being mothballed or shut down. This could be because only productive leases have many wells. The number of service wells is insignificant for mothballing and shutdown. The variables that condition the estimation on the two possible initial states shutdown and new entry are both statistically insignificant, which indicates that the correction for the initial problem is not important.

The variable *m(t-1)* is the lagged mothballing indicator. Based on the production, mothballing, and shutdown model in Chapter 4, a positive coefficient estimate on this variable indicates that there are mothballing sunk costs and that they play a role in the production decision. With the probit specification, the estimates for *m(t-1)* are .670 and 1.056 for the mothballing and shutdown state, respectively. This result implies that there are mothballing sunk costs.

In theory, the options value of an investment opportunity is increasing in sunk costs, implying that state dependence increases in sunk costs. I interact the lagged mothballing indicator *m(t-1)* with Depth to test this conjecture. The coefficients on Depth**m(t-1)* is .143 for the mothballing choice. Depth increases the probability of a mothballed lease to remain mothballed. However, this variable is insignificant for the shutdown choice. It does not affect the shutdown probability of leases that are mothballed.

Assuming that the sunk costs on an individual lease are increasing in the number of wells, the lagged number of wells on a lease is interacted with *m(t-1)*. As expected, the coefficient on this variable is positive and significant. The estimate is .091 for mothballing, and .088 for shutdown. The more wells there are on a lease, or the larger the sunk costs, the more the lagged mothballing status matters.

The final five rows of *Table 6–2* reveal how lagged shutdown, *s(t-1)*=1, affects current production, mothballing, and shutdown. The coefficients on *s(t-1)* and interacted terms are interpreted analogous to the lagged mothballing estimates. Because there are no transition from shutdown to mothballing, there are no estimates on *s(t-1)* in the mothballing columns.

The lagged shutdown status is significant. Its estimate is 2.924. A lease that is shutdown is very likely to remain shutdown and I conclude that there are shutdown sunk costs that matter for a lease's current production status. Note, that equation 6.5 implies that the coefficient estimate on *s(t-1)* should be larger than the one on *m(t-1)*. The empirical results confirm this conjecture. The sunk costs of shutting down and returning to production are larger than the ones of mothballing and returning to production.

Finally, consider Depth, and the lagged number of wells on a lease interacted with *s(t-1)*. The interaction of Depth and *s(t-1)* is positive and significant at the 10 percent level. There is some evidence that the sunk costs involved in shutting down are increasing in the depth of a well. However, the lagged number of wells is insignificant.

Overall there is strong evidence that sunk costs matter for the current production status of an oil lease. I find that both states, mothballing and shutdown, carry sunk costs. The sunk costs involved in shutdown appear to be larger. Moreover, the costs of mothballing are increasing in the average well depth on a lease, and higher oil price volatility lowers the probability of switching from mothballing to shutdown. There is some evidence that the costs of shutting down and re-entering production are increasing in the average well depth.

6.4 No-Production vs. Mothballing and Shutdown

The three-choice model of production, mothballing, and shutdown is considerably harder to estimate than the two-choice model that combines mothballing and shutdown into no-production. What are the additional insights that estimating the three-choice model provides?

The coefficient estimate on the lagged mothballing *m(t-1)* is considerably smaller than the one on the shutdown indicator *s(t-1)*. There is more persistence in the shutdown state than in the mothballing state. The coefficient estimate on lagged no-production *np(t-1)* in the two-choice model lies between the estimates for *m(t-1)* and *s(t-1)*. Industry policy recommendation and implications regarding the closing and opening of wells or leases may be very different depending on the level of data that is used. For example, based on the two-choice model it may look feasible to induce oil producers to re-open leases with a particular subsidy. However, when using the three-choice model one may learn that the proposed subsidy would not be sufficient to induce the re-opening of shutdown leases.

In addition, it is learned that some variables only affect the decision to shutdown and others affect only the decision to mothball. For example, deeper wells are less likely to mothball and older wells are more likely to shutdown. At the same time, the variable Depth does statistically not affect the decision to shutdown and Age does not affect the decision to mothball. However, in the two-choice model the variable Depth lowers the probability of no-production and Age increases the probability of no-production. Again, this can have implications for policy recommendations.

6.5 Goodness of Fit and Predicted Probabilities

The coefficient estimates for the two-choice model and the three-choice model can be used to predicted transition probabilities for a lease given its characteristics, lagged production status, and the macro environment. *Table 6–3* and *Table 6–4*, for example, compute the average of the individual transition probabilities for each year. The predicted average transition probabilities are compared to the actual observed transition rates to assess the overall fit of the model. The actual probabilities are the ones reported in *Table 5.7* and *Table 5.8*.

Table 6–3 compares the predicted average transition probabilities to the actual transition rates for the two-choice model with the states production and no-production for the years 1980 to 1997. The predicted probabilities are computed based on the MSM Probit estimates in *Table 6–1*. Columns 2 and 3 in *Table 6–3* present the predicted and actual probabilities, respectively, for switching from the production state to the no-production state. Similarly, columns 4 and 5 present the probabilities of remaining in the no-production state.

The lowest predicted and actual transition rates from production to no-production are .019 and .009, respectively. The highest predicted and actual transition probabilities for the same event are .071 and .088, respectively. The averages of the annual predicted probabilities and observed probabilities are .047 and .043. Generally, the model seems to capture the transition from production to no-production reasonably well.

The predicted and actual transition rates out of the no-productions are generally of the same magnitude. The lowest predicted and actual transition rates are .766 and .617, respectively. The highest predicted and actual rates are .876 and .927, respectively. In any case, movements out of the no-production state do not appear to be captured as well as the movements out of the production state. For example, there are several years, 1982, 1983, 1984, and 1988, for which the actual probabilities of staying in the no-production state are around .6 and the predicted probabilities are around .8. Overall, the predicted probabilities out of the no-production state do not fluctuate over time as much as the observed transition probabilities. At this point, it is not clear what drive these results. The three-choice model that distinguishes between mothballing and shutdown provides further insights.

Table 6–4 displays the annual average predicted and observed transition probabilities for the production, mothballing, and shutdown state. The predicted probabilities are based on the MSM estimates for the three-choice model in *Table 6–2*. The predicted and actual probabilities for switching from production to mothballing are displayed in columns 2 and 3, respectively.

Similarly, columns 4 and 5 present the predicted and actual probabilities for transition from production to shutdown. Columns 6 to 9 summarize the transition probabilities out of the mothballing state. The final two columns show the predicted and observed probabilities out of shutdown.[41]

Transitions out of production appear to be well explained by the model. For example, for the years 1989, 1990, and 1991 the predicted transition rates from production to mothballing are .027, .017, and .026, and the observed ones are .035 .014, and .034. For the same years the predicted transition rates out of production are .027, .014, and .026, and the observed ones are .022, .006, and .031. The model seems to be capable of picking up patterns that affect both non-producing states. However, in the years 1986 and 1987 the observed transition probabilities from production to mothballing are relatively high and the observed production to shutdown probabilities are very low. At the same time, the predicted probabilities for switching form production to mothballing or shutdown are both in the mid range of the predicted probabilities. In contrast,, in 1996 the observed transition rate from production to shutdown is at its highest, .076, and the observed transition rate from production to mothballing is relatively low. The model predicts somewhat above average transition probabilities for mothballing and shutdown in 1996. It appears that the model cannot pick up factors that affect only the value of one of the two non-producing states.

Columns 6 through 9 report the predicted and observed transition probabilities out of the mothballing state. The picture here is two fold. In most years, about 30 percent of the mothballed leases stay mothballed and this is the generally what the model predicts. However, there are several years with outliers in the actual probabilities that the model is not capable explaining: 1981 (.667), 1984 (.714), 1991 (.167), and 1993 (.190). Similarly, for transition from

[41] Recall, there are no transition from shutdown to mothballing.

mothballing to shutdown, the model predicts transition rates between .146 and .418, and the actual transition rates are between .000 and .385. Here, the model is not capable explaining some of the relatively low transition rates that look like outliers, for example 1981 (.000), 1984 (.071), 1992 (.056). On the same hand, in 1981, 1984, and 1992 the observed re-mothballing probabilities are extremely high with, .667, .714, and .528. Recall that there is a maximum legal time limit for wells to be mothballed. Transition from mothballing to production, mothballing, or shutdown are partly policy driven and the comparison of predicted and observed transition probabilities indicates that there were policy changes that the model cannot pick up.

The final two columns report the annual predicted and observed transition rates from shutdown to shutdown. The model predicts probabilities ranging from .750 to .920, while the actual transition rates are between .610 and .958. Although the model's predictions are generally of the same magnitude as the observed transition rate, it does not perform very well in predicting the very low actual transition rates in 1983 (.658), 1984 (.610), and 1988 (.652). For example, the predicted rate for 1988 is .905, which is .253 above the actual rate. Nevertheless, the very low probability of staying shutdown appears to be driven by something outside the scope of the model. All observed transition rates since 1986, except from the one for the year 1988, are larger than .800. In addition, the nominal oil price is at its lowest in 1988, and the real oil price at its second lowest in 1988, see *Table 5-1*. This is another indication that policies were in place that the model is not capable picking up.

6.5.1 Transition Probabilities and the Price of Oil

Table 6–5 predicts how oil producers would have reacted to different oil prices in 1997. The table reports the average predicted transition probabilities from 1996 to 1997 using the actual price of oil and fictitious lower and higher oil prices. The predicted probabilities are computed using the coefficient estimates for the two and three-choice models that are estimated with S.D. Price and Depth. The predicted transition probabilities are the choice probabilities conditional on the lagged production state, and, therefore, the predicted probabilities are computed using only the observations in the respective lagged production state.

The first and second row of *Table 6–5* report the actual and fictitious nominal and real prices in 1997. The actual nominal price in 1997 is $20.53/bbl. The model is estimated with the real price of oil, base year 1982, which is $16.09/bbl in 1997. The oil prices to the left are lower than the actual price and the prices to the right are higher than the actual price. Rows 4 and 5 plot the predicted average transition probabilities for the two-choice model. Rows 7 to 11 display the average predicted transition probabilities for the three-choice model.

Given the actual price of oil, the average predicted probability in 1997 for switches from production to no-production is .062, and the average predicted probability of remaining in the no-production state is .867. These transition rates change with the price of oil. *Table 6–5* computes average predicted transition probabilities for nominal oil prices between $7.66/bbl and $58.69/bbl. The probability of switching from production to no-production and the probability of remaining in the no-production state are both decreasing in the price of oil. The probabilities of switching from production to no-production range from .020 to .146, and the predicted probabilities for staying in the no-production state are between .724 and .944. The table implies that even at an oil price of $7.66/bbl only 14.6 percent of producing leases switch into the no-production state. Considering transition out of the no-production state, 76.3 percent of all leases remain in the no-production state even at the extraordinary high price of $45.39/bbl.

The lowest 5 rows of *Table 6–5* present the average predicted probabilities for the three-choice model. As in the two-choice model, the predicted transition probabilities are decreasing in the price of oil.[42] The predicted probabilities for transitions from production to mothballing are .050 and .013 at the lowest and highest fictitious oil price, respectively. The predicted probabilities for transitions from production are .109 and .006 at the lowest and highest fictitious oil price, respectively. First, these results show that the production state is very persistence and that even large price changes generate only small changes in state persistence. Second, these results indicate that transitions from production to shutdown are more price-sensitive than transitions from production to mothballing. Leases switch to mothballing at a broad range of prices, where else leases switch to shutdown mainly at very low oil prices.

The predicted transition probabilities from mothballing to shutdown or mothballing reveal a similar pattern. The predicted transition rates for switching from mothballing to shutdown are between .240 and .460. The predicted transition rates from mothballing to shutdown range from .120 to .536, and they are strongly impacted by the price of oil. Again, at the high price leases are more likely to switch to mothballing than to shutdown. At the low price the shutdown probability is higher than the mothballing probability. Transitions from mothballing to shutdown appear to be price-sensitive.

Finally, the last row of *Table 6–5* presents the predicted transition rates for staying in the shutdown state. The predicted probabilities range from .751 to .973, and they are decreasing in the price of oil. The model predicts that even at an oil price of $39.55/bbl, 82 percent of the shut down leases do not return to production. The shutdown state displays a lot of state persistence and even large price changes generate only little changes in the predicted lagged-shutdown transition probability.

---

[42] Note that the probability of being in the production state is inverse to the displayed probabilities and, thus, increasing in the price of oil.

6.5.2 Transition Probabilities and Latent Lease Values

Table 6–6 computes predicted transition probabilities by latent profit and lease value difference categories. The two-choice model estimates sunk costs, and the latent profit and value differences between production and no-production $\alpha + \boldsymbol{\beta} \boldsymbol{P}_t + \boldsymbol{\delta} \boldsymbol{X}_{it}$ (see equation 6.3). The three-choice model estimates sunk costs, and the latent profit and value differences of mothballing and production $\alpha_2 + \boldsymbol{\beta}_2 \boldsymbol{P}_t + \boldsymbol{\delta}_2 \boldsymbol{X}_{it}$, and shutdown and production $\alpha_3 + \boldsymbol{\beta}_3 \boldsymbol{P}_t + \boldsymbol{\delta}_3 \boldsymbol{X}_{it}$ (see equation 6.6). The larger these terms are the more valuable are the states no-production, mothballing, or shutdown compared to production. *Table 6–6* predicts these terms, given a lease's lagged production state and sorts them by size. Their size is increasing from left to right (see row 1). The first category contains the 5 percentile, the second category contains the 5 to 10-percentile, and so on. The predicted transition probability is computed for each lagged-production/size category cell, using the average predicted sunk costs for a particular transition out of the given lagged production state.

Consider the two-choice model first in row 3 and 4. The probability of switching from production to no-production and the probability of staying in the no-production state are both increasing in the sort, as the value of the no-production state is increasing in the size categories. The probability of switching from production to no-production ranges from .006 to .0112 and does not change very much across the latent profit and lease value percentiles. However, the probability of staying out of production is between .624 and .931, and it varies a lot across the sort. Latent profits and lease values seem to be more important for the entry than for the exit decision.

The lowest 5 rows display the predicted transition probabilities for the three-choice model. The predicted probabilities for transitions out of the production state are between .004 and .073 for mothballing, and .002 and .053 for shutdown. Even when the relative values of the states mothballing and shutdown are at their highest, the probability stopping production is quite low with .073 and .053. The production state is very persistent and leases are unlikely to mothball or shutdown.

There is much more variation in the transition rates from mothballing to mothballing or shutdown than there is variation out of the production state. The predicted probabilities of remaining in the mothballing state are between .109 and .488. The predicted probabilities for switching from mothballing to shutdown are similar and range from .057 to .382. The decision to stay mothballed or to switch from mothballing to shutdown is heavily affected by latent profits and lease values. This implies, that mothballed leases return to production if production seems profitable.

For leases that are shutdown, the most profitable ones return to production with a probability of .329. On the other hand, the least profitable ones remain shut down with .943 probability. The probability of remaining shutdown is only for leases in the 25 percentile smaller than .800. In generally, shutdown leases are not very likely to return to production.

6.5.3 Transition Probabilities and Sunk Costs

Table 6–7 computes transition probabilities by sunk cost percentiles. Sunk costs are based on the predicted values of $\boldsymbol{\gamma} np_{it-1}\boldsymbol{H}_{it}$ in the two-choice case (see equation 6.3) and $\boldsymbol{\gamma}_2 m_{it-1}\boldsymbol{H}_{\boldsymbol{it}}$ and $\boldsymbol{\gamma}_3 m_{it-1}\boldsymbol{H}_{\boldsymbol{it}}$ in the three-choice case (see equation 6.6). The predicted average probabilities are computed using all observation that fall into a given lagged production state/sunk cost category cell. The latent profit and lease values for an individual lease are replaced by the average for all leases in the same lagged state. The sunk costs categories in *Table 6–7* are increasing from left to right.

The third row displays the average predicted probability that a lease in the no-production state remains non-producing. As expected, the probability of remaining in the no-production state is increasing in sunk costs, namely from .685 to .954. Leases with the highest level of sunk costs are very unlikely to return to production.

The three-choice model estimates show that the transitions from mothballing to mothballing or shutdown are strongly affected by the level of sunk costs. For the lowest category of sunk costs the predicted transition rates are .155 and .208 for staying mothballed and switching to shutdown, respectively. The same probabilities for the highest sunk costs category are .666 and .569. Mothballed low-sunk costs leases are much more likely to return to production than high sunk costs leases.

Leases in the shutdown state are affected by different levels of sunk costs, too. The probabilities of remaining shutdown range from .846 to .927. Although the probability of staying shutdown increases in sunk costs, shutdown leases are not very likely to return to production regardless of their respective sunk costs category.

6.6 Summary

Based on an oil lease data set from Oklahoma, this chapter reports estimates for two dynamic discrete-choice models. The first distinguishes the production and no-production decision, while the second distinguishes among production, mothballing, and shutdown. For estimating the three-choice model, the data is refined such that the no-production state is distinguished as either mothballing or shutdown. As outlined in Chapter 4 it is possible to test for sunk costs by testing for state persistence. The lagged production status matters for the current production status if sunk costs matter. If there are no sunk costs involved in switching the production status, then the current production status is independent of the lagged production history because the firm can adjust its status each period at no costs. Furthermore, if there are no sunk costs, then variables that affect only future expected profits, for example uncertainty over future oil prices, should not affect the current production status.

Using aggregate data from Oklahoma and Texas, Chapter 3 provided first evidence that oil recovery is a sunk costs industry. As suggested by this macro evidence, I also find strong evidence of sunk costs in the micro-data from Oklahoma. The coefficient estimate on the lagged production indicator *np(t-1)* in the two-choice model, and the estimates on the lagged production state indicators *m(t-1)* and *s(t-1)* are positive and significant. Based on the theoretical model in Chapter 4, I conclude that this evidence of state dependence implies that oil production is a sunk costs industry. Sunk costs are an important determinant of the production decision, and are important in explaining the entry and exit hysteresis that is observed in the oil industry.

In addition, there is some evidence that the sunk mothballing and shutdown costs are increasing in the depth of a well. This is a sensible finding as workovers and re-drillings of deeper wells are more costly. Uncertainty over future market conditions, measured by the within-year oil price volatility, also affects the production state choice. This can be attributed to the options value of investment, which is increasing in uncertainty.

The probability of choosing a non-producing state is decreasing in the price of oil. A higher price of oil increases the probability of oil production. However, the price of oil is not as important for explaining production as the prior production status. Because of the importance of sunk costs, producing oil leases have a low probability of exiting and non-producing leases have a low probability of re-entering production.

Table 6–1: Estimates for the Two-Choice Model: Production and No-Production

Estimates for the Two-Choice Model: Production and No-Production

	Probit		MSM Probit	
	Coef.	T-stat	Coef.	T-stat
Intercept	0.100	0.532	-0.007	-0.035
Oil Price	-0.495	-9.195	-0.505	-9.587
S.D. Price	-0.968	-3.033	-0.980	-3.109
Oil Production	-0.007	-5.339	-0.009	-5.125
Depth	-0.093	-4.415	-0.083	-3.173
$Depth^2$	0.003	1.233	0.002	0.831
Age	0.160	4.711	0.233	5.713
Age^2	-0.002	-4.289	-0.003	-4.820
# of Wells	-0.051	-6.228	-0.052	-5.457
# of Service Wells	0.046	1.970	0.048	1.715
I(Initially No-Production)	0.034	0.900	0.099	2.057
np(t-1)	1.884	21.667	1.916	18.679
Depth**np(t-1)*	0.130	9.070	0.131	8.696
# of Wells**np(t-1)*	0.054	4.555	0.049	3.938
# of past *np***np(t-1)*	0.036	3.838	-0.021	-1.939
λ		-	0.136	5.403
ρ		-	-0.244	-8.472

The choice production is the base category. The table estimates the probability of being in the state no-production. Number of observations=18306.

I(Initially No-Production) = 1 if the first observation for a lease is no-production, 0 otherwise (initial condition correction); *np(t-1)*= 1 if the lease is in the no-production state at *t*-1, 0 otherwise; # of past *np* = number of consecutive periods of no-production.

Table 6–2: Estimates for the Three-Choice Model: Production, Mothballing, and Shutdown

Estimates for the Three-Choice Model: Production, Mothballing, and Shutdown								
	Logit				MSM Probit			
	M		**S**		**M**		**S**	
	Coef.	T-stat	Coef.	T-stat	Coef.	T-stat	Coef.	T-stat
Intercept	-0.118	-0.188	-0.148	-0.252	-0.307	-1.032	-0.396	-1.415
Oil Price	-0.607	-3.602	-1.363	-8.642	-0.300	-3.947	-0.634	-8.797
S.D. Price	-1.702	-1.830	-1.592	-1.922	-0.813	-1.909	-0.887	-2.306
Oil Production	-0.014	-2.498	-0.019	-4.615	-0.004	-1.848	-0.009	-4.209
Depth	-0.418	-6.326	0.013	0.217	-0.210	-6.703	0.014	0.451
$Depth^2$	0.017	2.284	-0.002	-0.405	0.011	3.347	-0.002	-0.658
Age	0.090	0.773	0.482	4.399	0.061	1.040	0.282	4.889
Age^2	-0.002	-1.254	-0.006	-4.000	-0.001	-1.537	-0.003	-4.335
# of Wells	-0.126	-5.163	-0.093	-3.329	-0.051	-4.567	-0.044	-3.546
# of Service W.	0.099	1.387	0.104	1.622	0.049	1.386	0.054	1.622
I(Initially Shut)	0.142	0.915	-0.253	-2.048	0.074	0.930	-0.032	-0.472
I(Initially New)	-0.098	-0.536	0.113	0.664	-0.012	-0.123	0.095	1.068
m(t-1)	1.591	6.234	2.332	7.379	0.670	4.712	1.056	6.380
Depth**m(t-1)*	0.356	7.520	0.097	1.713	0.143	5.582	0.052	1.743
# of Wells**m(t-1)*	0.222	4.837	0.201	3.943	0.091	3.853	0.088	3.320
s(t-1)	-		5.838	24.522	-		2.924	23.722
Depth**s(t-1)*	-		0.062	1.592	-		0.037	1.917
# of Wells**s(t-1)*	-		0.061	1.902	-		0.024	1.543
		Coef.	T-stat			Coef.	T-stat	
λ			-			0.114	6.087	
ρ_1			-			0.000	fixed	
ρ_2			-			0.000	fixed	
ϕ_{12}			-			0.000	fixed	
ϕ_{22}			-			1.000	fixed	
ψ_{12}			-			0.000	fixed	
ψ_{22}			-			1.000	fixed	

The choice Production is the base category. The table estimates the probabilities of the states mothballing and shutdown. Number of observations=18306.

M=Mothballing; S=Shutdown; I(Initially Shut) = 1 if the first observation for a lease is shutdown, 0 otherwise (initial condition correction); I(Initially New) = 1 if a lease enters during the data period, 0 otherwise (initial condition correction); *m(t-1)* = 1 if mothballed at *t-1*, 0 otherwise; *s(t-1)* =1 if shutdown at *t-1*, 0 otherwise.

Table 6–3: Goodness of Fit, Two Choices: Predicted and Observed Probabilities

Goodness of Fit, Two Choices: Predicted and Observed Probabilities				
State *t*-1	Production		No-Production	
State *t*	No-Production		No-Production	
Year *t*	Predicted	Observed	Predicted	Observed
1980	0.019	0.020	0.766	0.669
1981	0.023	0.009	0.803	0.738
1982	0.025	0.015	0.808	0.679
1983	0.031	0.043	0.824	0.653
1984	0.032	0.030	0.812	0.647
1985	0.034	0.023	0.806	0.707
1986	0.051	0.040	0.843	0.803
1987	0.050	0.038	0.841	0.755
1988	0.061	0.044	0.852	0.617
1989	0.056	0.057	0.838	0.716
1990	0.031	0.019	0.774	0.771
1991	0.053	0.065	0.838	0.848
1992	0.057	0.048	0.853	0.834
1993	0.061	0.057	0.858	0.829
1994	0.067	0.080	0.875	0.854
1995	0.071	0.054	0.876	0.906
1996	0.055	0.088	0.851	0.869
1997	0.062	0.050	0.867	0.927
Mean	0.047	0.043	0.832	0.768

The predicted probabilities are based on the MSM Probit estimates in Table 6.1.

Table 6–4: Goodness of Fit, Three Choices: Predicted and Observed Transition Probabilities

Goodness of Fit, Three Choices: Predicted and Observed Transition Probabilities

	State *t-1*	Production				Mothballing				Shutdown	
	State *t*	Mothballing		Shutdown		Mothballing		Shutdown		Shutdown	
Year *t*		Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed
1980		0.015	0.019	0.006	0.001	-	-	-	-	0.750	0.669
1981		0.017	0.009	0.008	0.000	0.264	0.667	0.135	0.000	0.800	0.746
1982		0.018	0.011	0.009	0.004	0.266	0.235	0.137	0.294	0.814	0.702
1983		0.020	0.028	0.011	0.016	0.285	0.385	0.161	0.231	0.835	0.658
1984		0.020	0.023	0.012	0.007	0.297	0.714	0.189	0.071	0.834	0.610
1985		0.021	0.017	0.013	0.006	0.306	0.308	0.204	0.385	0.839	0.713
1986		0.025	0.031	0.026	0.009	0.334	0.370	0.283	0.370	0.887	0.821
1987		0.026	0.032	0.024	0.006	0.337	0.378	0.240	0.189	0.887	0.818
1988		0.028	0.028	0.032	0.016	0.357	0.357	0.282	0.167	0.905	0.652
1989		0.027	0.035	0.027	0.022	0.334	0.342	0.274	0.184	0.891	0.782
1990		0.017	0.014	0.014	0.006	0.278	0.409	0.220	0.295	0.828	0.794
1991		0.026	0.034	0.026	0.031	0.361	0.167	0.286	0.300	0.886	0.938
1992		0.027	0.026	0.028	0.022	0.350	0.528	0.293	0.056	0.896	0.890
1993		0.027	0.026	0.032	0.031	0.340	0.190	0.289	0.381	0.905	0.893
1994		0.029	0.033	0.036	0.047	0.373	0.467	0.317	0.133	0.917	0.893
1995		0.030	0.021	0.039	0.033	0.356	0.390	0.325	0.244	0.920	0.956
1996		0.024	0.013	0.028	0.076	0.330	0.242	0.254	0.273	0.898	0.913
1997		0.027	0.017	0.032	0.033	0.347	0.278	0.301	0.111	0.908	0.958
Mean		0.023	0.023	0.022	0.020	0.324	0.378	0.246	0.217	0.866	0.800

The predicted probabilities are based on the MSM Probit estimates in Table 6.2.

Table 6–5: Predicted Transition Probabilities for 1996/1997 with Observed Oil Price and Fictitious Oil Prices

Predicted Transition Probabilities for 1996/1997 with Observed Oil Price and Fictitious Oil Prices											
	Nominal Oil Price	$7.66	$14.04	$17.86	**$20.53***	$22.97	$26.79	$33.17	$39.55	$45.93	$58.69
	Real Oil Price (Used in Estimation)	$6.00	$11.00	$14.00	**$16.09 ***	$18.00	$21.00	$26.00	$31.00	$36.00	$46.00
State in 1996	State in 1997	Two-Choice Model: Production and No-Production									
Production	No-Production	0.146	0.088	0.071	**0.062**	0.055	0.047	0.038	0.031	0.026	0.020
No-Production	No-Production	0.944	0.903	0.881	**0.867**	0.855	0.837	0.809	0.785	0.763	0.724
		Three-Choice Model: Production, Mothballing, and Shutdown									
Production	Mothballing	0.050	0.034	0.029	**0.027**	0.025	0.022	0.019	0.017	0.015	0.013
	Shutdown	0.109	0.053	0.039	**0.032**	0.027	0.022	0.016	0.012	0.009	0.006
Mothballing	Mothballing	0.460	0.389	0.362	**0.347**	0.335	0.318	0.296	0.278	0.264	0.240
	Shutdown	0.536	0.387	0.331	**0.301**	0.277	0.246	0.206	0.176	0.153	0.120
Shutdown	Shutdown	0.973	0.940	0.921	**0.908**	0.896	0.877	0.848	0.822	0.796	0.751

* The price in bold font represents the actual average oil price in 1997.
The predicted probabilities for the two-choice model are based on the MSM Probit estimates in Table 6.1.
The predicted probabilities for the three-choice model are based on the MSM Probit estimates in Table 6.2.

Table 6–6: Predicted Transition Probabilities by Latent Profit and Lease Value Difference Percentiles

Predicted Transition Probabilities by Latent Profit and Lease Value Difference Percentiles									
	Percentile	0-.05	.05-.10	.10-.25	.25-.50	.50-.75	.75-.90	90-.95	.95-1.00
State *t*-1	State *t*	Two-Choice Model: Production and No-Production							
Production	No-Production	0.006	0.012	0.021	0.035	0.053	0.073	0.089	0.112
No-Production	No-Production	0.624	0.735	0.785	0.833	0.873	0.897	0.913	0.931
		Three-Choice Model: Production, Mothballing, and Shutdown							
Production	Mothballing	0.004	0.007	0.010	0.014	0.022	0.039	0.058	0.073
	Shutdown	0.002	0.004	0.008	0.015	0.026	0.038	0.045	0.053
Mothballing	Mothballing	0.109	0.152	0.182	0.220	0.275	0.367	0.440	0.488
	Shutdown	0.057	0.098	0.137	0.199	0.265	0.324	0.353	0.382
Shutdown	Shutdown	0.671	0.774	0.822	0.879	0.912	0.928	0.936	0.943

The predicted probabilities for the two-choice model are based on the MSM Probit estimates in Table 6.1.
The predicted probabilities for the three-choice model are based on the MSM Probit estimates in Table 6.2.

Table 6–7 Predicted Transition Probabilities by Sunk Cost Percentiles

Predicted Transition Probabilities by Sunk Cost Percentiles									
	Percentile	0-.05	.05-.10	.10-.25	.25-.50	.50-.75	.75-.90	90-.95	.95-1.00
State *t*-1	State *t*	Two-Choice Model: Production and No-Production							
No-Production	No-Production	0.685	0.706	0.747	0.828	0.871	0.905	0.926	0.954
		Three-Choice Model: Production, Mothballing, and Shutdown							
Mothballed	Mothballed	0.155	0.173	0.198	0.265	0.369	0.474	0.556	0.666
	Shutdown	0.208	0.209	0.209	0.216	0.252	0.319	0.405	0.569
Shutdown	Shutdown	0.846	0.851	0.866	0.879	0.889	0.899	0.906	0.927

The predicted probabilities for the two-choice model are based on the MSM Probit estimates in Table 6.1.
The predicted probabilities for the three-choice model are based on the MSM Probit estimates in Table 6.2.

Chapter 7

SUMMARY AND CONCLUSIONS

This chapter summarizes the findings of this dissertation. Based on the empirical findings, I draw conclusions and make policy recommendations. Finally, suggestions for future research are made.

7.1 Summary and Conclusions

This dissertation analyzes the production, mothballing, and shutdown behavior of oil producers in Oklahoma from 1980 to 1997. Based on theoretical findings, the hypothesis is that sunk costs play an important role for production decisions in the oil industry. Theoretical models that analyze investment and production when costs are sunk and future market condition are uncertain, find that the combination of uncertainty and sunk costs generate entry and exit hysteresis, and hysteresis in switching among production states.

Aggregate data from Oklahoma and Texas provide first evidence that sunk cost models apply to the oil industry. The second oil price shock, with the real price of oil at an unmatched high, induced the drilling of many new oil wells and increased the total number of wells in Oklahoma and Texas. However, when the real price fell back to its initial level the total number of wells did not fall back to its pre-shock level. Most of the new wells stayed in operation, although the effect that had induced their opening had disappeared. This is a clear picture of entry and exit hysteresis.

Nevertheless, the question of how an individual oil production unit reacts to changes in the price of oil remains unclear. To investigate the production behavior of individual oil producing units, I construct a new data set for oil leases in Oklahoma. The decision to produce, mothball, or shutdown an oil lease is modeled using a dynamic program. Switching the production status of the lease is costly. If costs of switching the production status are sunk, the lagged production status matters for the current production choice. Moreover, if there are no sunk costs, future expected lease value does not depend on the current production status and variables that affect only a lease's future expected value should not matter for the current production status.

The dynamic program's Bellman equation can be formulated as an empirical discrete choice model and estimated with the logit or the probit estimator. By including lagged dependent state indicators into the regression, it is possible to test for the absence of sunk costs. Interacting the lagged dependent state indicator with lease characteristics allows sunk costs to vary across leases. In addition, the model allows investigation of how variables that determine an oil lease's productivity affect the probabilities of being in each production state.

As expected, it is found that the lagged production status matters. I conclude that sunk costs play an important role in the production decision. Moreover, state persistence is increasing in the well depth, a proxy for sunk costs. Providing further evidence that sunk costs model are relevant, the within-year oil price volatility, a measure of uncertainty over future market conditions, affects the current production status, too. It is found that the probability of being in production is increasing in the price of oil. Nevertheless, an oil lease's current production status is mainly affected by it production status in the previous year.

The model with the states mothballing and shutdown is empirically distinct from the simpler model that just distinguishes between production and no-production. First, the sunk costs involved in shutdown appear to be larger than the sunk costs involved in mothballing. Second, some variables are only significant for the mothballing decision and others are only significant for the shutdown decision. In the two-choice model on the other hand, no-production is affected by all these variables. Thus, conclusions based on the empirical results are different depending on the level of data that is used.

7.2 Policy Recommendations

Politicians in oil producing states with high variable costs stripper wells, for example Oklahoma and Texas, are interested in prolonging the production life of their depleted oil fields. The sunk costs involved in shutting down seem to be very substantial. Attempting to induce leases and wells to return to production after being shutdown would probably be very costly. In fact, the oil industry asserts that remaining reserves are lost forever once these low productivity wells are shutdown. Based on my empirical results this seems to be true and it seems to be important to prevent leases from shutting down in the first place.

The empirical results show that shutdown is not only driven by the price of oil, but also by well and lease characteristics, such as the age and production rates. Considering industry policies, it seems to be important to investigate further as to which leases should be kept in production. It may be possible to help wells to stay open that are forced to close because of a temporarily low oil price, but not the ones that are about to shutdown even at a reasonable high price of oil.

Many leases that mothball eventually will shutdown. It has already been recognized that it may be very costly to enforce shutdown after a relatively short period of mothballing. Actually, selected wells are granted the permission to be mothballed for more than a year. An interesting possibility of preventing shutdown is to lengthen the mothballing period. This policy can be studied with data that contains indicators on which leases were exempt from the general plugging and abandonment requirements, and information on the replacing abandonment rule.

7.3 Suggestions for Future Research

This study can only make policy recommendations based on the status quo, but it is not suited for studying and predicting how the oil industry reacts towards policy changes. However, industry behavior can be simulated with a structurally estimated model. Thus, it appears to be important to estimate structural form of the dynamic program of oil production in Chapter 4.

BIBLIOGRAPHY

Abel, Andrew B. and Janice C. Eberly (1994). "A Unified Model of Investment Under Uncertainty." *The American Economic Review* 84(5): 1369-1384.

American Petroleum Institute (1999a), *Basic Petroleum Data Book: Petroleum Industry Statistics*. Volume XIX, Number1 (January). Washington D.C.: American Petroleum Institute.

American Petroleum Institute (1999b). "Economic State of the US Oil and Natural Gas Exploration and Production Industry: Long Term Trends and Recent Events." [Online]. Available at: http://www.api.org/pasp/upstream6.pdf (March 2, 2001).

Arrow, Kenneth J., and Anthony C. Fisher (1974). "Environmental Preservation, Uncertainty, and Irreversibility." *Quarterly Journal of Economics* 88(2): 312-319.

Baldwin, Richard, and Paul Krugman (1989). "Persistent Trade Effects of Large Exchange Rate Shocks." *Quarterly Journal of Economics* 104(4): 635-654.

Baldwin, Richard (1988). "Hysteresis in Import Prices: The Beachhead Effect." *The American Economic Review* 78(4): 773-785.

Bernanke, Ben S. (1983). "Irreversibility, Uncertainty, and Cyclical Investment." *Quarterly Journal of Economics* 98(1): 85-106.

Bresnahan, Timothy F. and Peter C. Reiss (1991). "Entry and Competition in Concentrated Markets." *Journal of Political Economy* 99(5): 977-1009.

Bresnahan, Timothy F., and Peter C. Reiss (1994). "Measuring the Importance of Sunk Costs." *Annales d'Economie et de Statistique* 34 (April-June): 181-217.

Campa, Jose M. (1994). "Entry by Foreign Firms in the United States under Exchange Rate Uncertainty." *Review of Economics and Statistics* 75(4): 614-622.

Campa, Jose M. (July 1998). "Exchange Rates and Trade: How Important Is Hysteresis in Trade?" *Manuscript*.

Chirinko, Robert (1993). **"**Business Fixed Investment Spending: Modeling Strategies, Empirical Results, and Policy Implications." *Journal of Economic Literature* 31(4): 1875-1911.

Cooper, Russel, John Haltiwanger, and Laura Power (1999). "Machine Replacement and the Business Cycle: Lumps and Bumps." *The American Economic Review* 89(4): 921-946.

Das, Sanghamitra (1992). "A Micro-Econometric Model of Capital Utilization and Retirement: The Case of the U.S. Cement Industry." *Review of Economic Studies* 59(2): 277-297.

Das, Sanghamitra, Mark Roberts, and James Tybout (April 2000). "Micro Foundations of Export Dynamics." *Manuscript*.

Department of Energy Information Administration (1995). " Oil and Gas Development in the United States in the Early 1990's: An Expanded Role for Independent Producers." [Online]. Available at: http://tonto.eia.doe.gov/FTPROOT/petroleum/0600.pdf (March 2, 2001).

Department of Energy Information Administration (1997). "Petroleum 1996 – Issues and Trends." [Online]. Available at:

ftp://ftp.eia.doe.gov/pub/oil_gas/petroleum/analysis_publications/petroleum_issues_trends_1996/ENTIRE.PDF (March 2, 2001).

Department of Energy Information Administration (1999). "Petroleum: An Energy Profile 1999." [Online]. Available at: ftp://ftp.eia.doe.gov/pub/oil_gas/petroleum/analysis_publications/petroleum_profile_1999/profile99v8.pdf (March 2, 2001).

Department of Energy Information Administration (2000). "U.S. Crude Oil, Natural Gas, and Natural Gas Liquids Reserves – 1999 Annual Report." [Online]. Available at: ftp://ftp.eia.doe.gov/pub/oil_gas/natural_gas/data_publications/crude_oil_natural_gas_reserves/historical/1999/pdf/arr99.pdf (March 2, 2001).

Department of Energy Information Administration (2001). "World Oil Market and Oil Price Chronologies: 1970 – 2000." [Online]. Available at: http://www.eia.doe.gov/emeu/cabs/chron.html (Feb. 28, 2001).

Dixit, Avinash (1989). "Entry and Exit Decisions under Uncertainty." *Journal of Political Economy* 97 (3): 620-638.

Dixit, Avinash (1992). "Investment and Hysteresis." *Journal of Economic Perspectives* 6(1): 107-132.

Doms, Mark, and Timothy Dunne (1998). "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics*: 409-429.

Dunne, Timothy, Mark J. Roberts, and Larry Samuelson (1988). "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries." *Rand Journal of Economics* 19(4): 495-515.

Dunne, Timothy, Mark J. Roberts, and Larry Samuelson (1989). "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics* 104(4): 671-698.

Garlick, David M., and Patricia Cleary Lee (1999). "Investments in Energy Security – State Incentives to Maximize Oil and Gas Recovery." [Online]. Available at: http://www.iogcc.state.ok.us/PDFS/Econ_Incentives_99_00.pdf (Feb. 28, 2001).

Green, William H. (2000), *Econometric Analysis.* New Jersey: Prentic Hall.

Hajivassiliou, V.A. (1993). "Simulation Estimation Methods for Limited Dependent Variable Models." In G.S. Maddala, C. Rao and H. D. Vinod, eds. *Handbook of Statistic* 11. Elsevier Science Publisher B.V., 1993.

Heckman, James J. (1981). "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process" in Charles Manski and Daniel McFadden, eds. *The Structural Analysis of Discrete Data.* Cambridge, MA: MIT Press: 179-195.

Henry, Claude (1974). "Investment Decisions Under Uncertainty: The Irreversibility Effect." *The American Economic Review* 64(6): 1006-1012.

Hubbard, Glenn R. (1994). "Investment Under Uncertainty: Keeping One's Option Open." *Journal of Economic Literature* 32(4): 1816-1831.

Hurn, A. S., and Robert E. Wright (1994). "Geology or Economics? Testing Models of Irreversible Investment Using North Sea Oil Data." *The Economic Journal* 104 (March): 363-371.

Interstate Oil and Gas Compact Commission (1999). “A Battle for Survival? The Real Story Behind Low Oil Prices.” [Online]. Available at http://www.iogcc.state.ok.us/PDFS/4-99lowpricereport.pdf (Feb. 28, 2001).

Interstate Oil and Gas Compact Commission (1999). “Assessing the Costs of Environmental Compliance.” [Online]. Available at: http://www.iogcc.state.ok.us/PDFS/Environmental%20Costs.pdf (March 2, 2001).

Interstate Oil and Gas Compact Commission (1999). “Marginal Oil and Gas: Fuel for Economic Growth.” [Online]. Available at: http://www.iogcc.state.ok.us/PDFS/99_Marg_O&G.pdf (March 2, 2001).

Keane, Michael P. (1994). "A Computationally Practical Simulation Estimator for Panel Data." *Econometrica* 62(1): 95-116.

Keane, Michael P. (1992). “A Note on Identification in the Multinomial Probit Model.” *Journal of Business & Economic Statistics* 10(2): 193-200.

Keane, Michael P. (1993). “Simulation Estimation for Panel Data Models with Limited Dependent Variables.” In G.S. Maddala, C. Rao and H. D. Vinod, eds. *Handbook of Statistic* 11. Elsevier Science Publisher B.V., 1993.

GEO Information Systems at Sarkeys Energy Center, University of Oklahoma. Available at: http://www.geo.ou.edu/nris/index.htm (May 22, 2001).

Geweke, John F., Michael P. Keane, David E. Runkle (1997). “Statistical Inference in the Multinomial Multiperiod Probit Model.” *Journal of Econometrics* 80(1): 125-165.

Lambson, Val Eugene, and Farell E. Jensen (1995). "Sunk Costs and the Variability of Firm Value over Time." *The Review of Economics and Statistics* 77(3): 535-544.

Lambson, Val Eugene, and Farell E. Jensen (1998). "Sunk Costs and Firm Value Variability: Theory and Evidence." *The American Economic Review* 88(1): 307-312.

Lowe, John S. (1982). *Fundamentals of Oil and Gas Leasing*. Tulsa, Oklahoma: IED Exploration Institute.

McDonald, Robert, and Daniel Siegel (1986). "The Value of Waiting to Invest." *The Quarterly Journal of Economics* 101(4): 707-728.

McFadden, Daniel (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica* 57(3): 995-1026.

National Petroleum Council (1994), *Marginal Wells.* National Petroleum Council.

Oklahoma Corporate Commission, Commission Rules. [Online]. Available at: http://www.occ.state.ok.us/TEXT_FILES/Rulesfrm.htm (Feb. 28, 2001).

Oklahoma Corporate Commission, Oil and Gas Information. [Online]. Available at: http://www.occ.state.ok.us/text_files/o&gfiles.htm (Feb. 28, 2001).

Oklahoma Marginal Well Commission (date unknown). "A Demographic and Economic Profile of Oklahoma's Marginal Oil and Gas Well Operators, Executive Summary." [Online]. Available at: http://www.state.ok.us/~marginal/Reports/Op_profile/op_profile.htm (Feb. 28, 2001).

Paddock, James L.; Daniel R. Siegel and James L. Smith (1988). "Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases." *Quarterly Journal of Economics* 103(3): 479-508.

Penn, David A., and John McCraw (date unknown). "Localized Impacts of Oil and Gas Production and Drilling Activity in Oklahoma of Oil and Gas Production and Drilling Activity in Oklahoma for the Oklahoma Commission on Marginally Producing Oil and Gas Wells." [Online]. Available at: http://www.state.ok.us/~marginal/localizedimpacts.htm (Feb. 28, 2001).

Penn, David A., and John McCraw (1996). "The Impact of Oil and Gas Production and Drilling on the Oklahoma Economy for Commission on Marginally Producing Oil and Gas Wells." [Online]. Available at: http://www.state.ok.us/~marginal/impactonok.htm (Feb. 28, 2001).

David A. Penn (1996). "Survey of Oklahoma Oil and Gas Leases for the Oklahoma Commission on Marginally Producing Oil and Gas Wells." [Online]. Available at: http://www.state.ok.us/~marginal/Surveyofleases.htm (Feb. 28, 2001).

Penn, David A. (1996). "Well Cost Analysis for the Oklahoma Commission on Marginally Producing Oil and Gas Wells." [Online]. Available at: http://www.state.ok.us/~marginal/wellcost.htm (Feb. 28, 2001).

Pindyck, Robert S. (1988). "Irreversible Investment, Capacity Choice, and the Value of the Firm." *American Economic Review* 78(5): 969-985.

Pindyck, Robert S. (1991). "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* 29(3): 1110-1148.

Roberts, Mark J. and James R. Tybout (1997). "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Cost." *The American Economic Review* 87(4): 545-564.

Stevenson, Paul, ed. (1988), *Oil and Gas Dictionary*. Great Britain: The Macmillan Press Ltd.

Wooldridge, Jeffrey M. (July 2000). "The Initial Conditions Problem in Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity." *Manuscript.*

WTRG Economists (1996-1999). "Oil Price History and Analysis." [Online]. Available at: http://www.wtrg.com/prices.htm (March 2, 2001).

WTRG Economists (1999). "Rotary and Workover Rig Counts." [Online]. Available at: http://www.wtrg.com/rotaryrigs.html (March 2, 2001).

VITA

Boris Molls

Education: Ph.D., Economics, The Pennsylvania State University

M.A., Economics, August 2000, The Pennsylvania State University

Vordiplom, Economics, February 1995, Rheinische Friedrich-Wilhelms Universitat Bonn, Germany

Areas of Concentration:

Primary: Industrial Organization, Econometrics

Secondary: Finance and Financial Econometrics

Relevant Professional Experience:

Research Assistant:

Prof. Mark J. Roberts (Summer 1999 to Fall 2001)

Teaching and General Education Assistant:

Statistics (Winter 1996, Universitat Bonn, Germany)

Introductory Statistics, Introductory Microeconomics, Introductory Macroeconomics, and Business Cycles (Spring 1997 to Spring 1999, Spring 2001, The Pennsylvania State University)

Professional Membership:

American Economic Association