Strategic Decision-Making with Information and Extraction Externalities: A Structural Model of the Multi-Stage Investment Timing Game in Offshore Petroleum Production

C.-Y. Cynthia Lin[1]

Abstract

When individual petroleum-producing firms make their exploration and development investment timing decisions, positive information externalities and negative extraction externalities may lead them to interact strategically with their neighbors. If they do occur, strategic interactions in petroleum production would lead to a loss in both firm profit and government royalty revenue. The possibility of strategic interactions thus poses a concern to policy-makers and affects the optimal government policy. This paper examines whether these inefficient strategic interactions take place on U.S. federal lands in the Gulf of Mexico. In particular, it analyzes whether a firm's production decisions and profits depend on the decisions of firms owning neighboring tracts of land. The empirical approach is to estimate a structural econometric model of the firms' multi-stage investment timing game. According to the results, when the tract sizes are large, firms do not impose externalities on each other on net when choosing to explore or develop, and, as a consequence, strategic considerations are second-order. This is the case with most of the tracts in the federal leasing program. However, in the few cases where the tract size is small, externalities do matter, and they cause firms to interact strategically with their neighbors. For small tracts, the effect of having a neighboring tract explored reduces real profits by about 26 million dollars, while having a neighboring tract developed raises real profits by about 3.5 million dollars. On the small tracts, the externalities cost about 17 million real dollars per tract developed.

Keywords: investment timing game, structural econometrics, petroleum production
JEL Classification: C21, C51, D92, L71
This draft: June 13, 2012

---

[1] University of California at Davis; cclin@primal.ucdavis.edu. I am indebted to William Hogan, Dale Jorgenson and Robert Stavins for their support and encouragement throughout this project. This paper also benefited from discussions with Susan Athey, Lori Bennear, Gary Chamberlain, Michael Greenstone, Joseph Harrington, Oliver Hart, Omowumi Iledare, Alan Krupnick, David Laibson, Adrian Lajous (President, Petrometrica), Charles Mason, Markus Mobius, Juan-Pablo Montero, Marcelo Moreira, Julie Mortimer, Michael Ostrovsky, Ariel Pakes, Jack Porter, Juan Robalino, Stephen Ryan, Michael Springborn, James Stock, Tomislav Vukina, Stephen Weinberg, and Martin Weitzman, among numerous others. I received helpful comments from participants at the First Bonzenfreies Colloquium on Market Dynamics and Quantitative Economics; the 2005 Harvard Environmental Economics Workshop; the 2006 International Industrial Organization Conference; the 3rd World Congress of Environmental and Resource Economics; the 9th Occasional Workshop on Environmental and Resource Economics at Santa Barbara; the 2008 AEA meetings; and at seminars at Harvard University, UC-Davis, UC-Berkeley, University of Hawaii at Manoa, Rice University, Duke University, University of North Carolina, UC-Riverside, and University of Alberta. I thank Kenneth Hendricks and Robert Porter for sharing their lease sale data with me. Patricia Fiore, April Larson and John Shaw helped me to understand the geology of oil production. Robert Yantosca provided assistance with the IDL software. I thank the Minerals Management Service, and especially John Rodi, Marshall Rose and Robert Zainey, for answering my many questions about the OCS leasing program. I am indebted to Bijan Mossavar-Rahmani (Chairman, Mondoil Corporation) and William Hogan for arranging for me to visit Apache Corporation's headquarters in Houston and a drilling rig and production platform offshore of Louisiana, and for their support of my research, and I thank the Repsol YPF - Harvard Kennedy School Energy Fellows Program for providing travel funds. I received financial support from an EPA Science to Achieve Results graduate fellowship, a National Science Foundation graduate research fellowship, a Repsol YPF - Harvard Kennedy School Pre-Doctoral Fellowship in energy policy, and from the Giannini Foundation of Agricultural Economics. All errors are my own.

1 Introduction

Petroleum production is a multi-stage process involving sequential investment decisions. The first stage is exploration: when a firm acquires a previously unexplored tract of land, it must first decide whether and when to invest in the drilling rigs needed to begin exploratory drilling. The second stage is development: after exploration has taken place, a firm must subsequently decide whether and when to invest in the production platforms needed to develop and extract the reserve. Because the profits from petroleum production depend on market conditions such as the oil price that vary stochastically over time, an individual firm producing in isolation that hopes to make dynamically optimal decisions would need to account for the option value to waiting before making either irreversible investment (Dixit & Pindyck, 1994).

The dynamic decision-making problem faced by a petroleum-producing firm is even more complicated when its profits are affected not only by exogenous market conditions, but also by the actions of other firms producing nearby. When firms own leases to neighboring tracts of land that may be located over a common pool of reserve, there are two types of externalities that add a strategic (or non-cooperative)[2] dimension to firms' investment timing decisions and may render these decisions socially inefficient.[3]

The first type of externality is an *information externality*: if tracts are located over a common pool or share common geological features so that their ex post values are correlated, then firms learn information about their own tracts when other firms drill exploratory wells on neighboring tracts (Hendricks & Porter, 1996). The information externality is a positive one, since a firm benefits from its neighbors' information. The information externality is socially inefficient because it may cause firms to play a non-cooperative timing game that leads them to inefficiently delay production, since the possibility of acquiring information from other firms may further enhance the option value to waiting. If firms are subject to a lease term by the end of which they must begin exploratory drilling, or else relinquish their lease, then the information externality would result in too little exploration at the beginning of the lease term and duplicative drilling in the final period of the lease (Hendricks &

[2] In this paper, I use the terms "strategic" and "non-cooperative" interchangeably.

[3] In my broad definition of an externality, I say that an externality is present whenever a non-coordinated decision by individual firms is not socially optimal.

Porter, 1996; Porter, 1995). Theoretical and empirical evidence for the information externality during exploration have been presented previously in a seminal series of papers by Kenneth Hendricks, Robert Porter and their co-authors (see e.g. Hendricks & Kovenock, 1989; Hendricks & Porter, 1993; Hendricks & Porter, 1996).

A second type of externality is an *extraction externality*: when firms have competing rights to a common-pool resource, strategic considerations may lead them to extract at an inefficiently high rate (Libecap & Smith, 2001; Libecap & Wiggins, 1985). When oil is extracted too quickly, it may cause a collapse in the formation being extracted from, thus collapsing the pipe and decreasing the total amount of oil extracted (Chermak et al., 1999). The extraction externality is a negative one, since it induces a firm to produce inefficiently. Libecap and Smith (2001) provide theoretical evidence for the extraction externality.

Owing to both information and extraction externalities, the dynamic decision-making problem faced by a petroleum-producing firm is not merely a single-agent problem, but rather can be viewed as a multi-agent, non-cooperative game in which firms behave strategically and base their exploration and development policies on those of their neighbors.

Both externalities lead to strategic interactions that are socially inefficient. The information externality leads to an inefficient delay in exploration. The extraction externality leads to excessively high extraction rates and less total oil extracted. Both types of strategic behavior lead to lower profits for the petroleum-producing firms and lower royalty revenue for the federal government. It is therefore important to analyze whether these strategic interactions place, and therefore whether policies that can mitigate the strategic interactions should be implemented. These policies include modifying the tract size, enacting policies to facilitate joint ventures and cooperation, and selling non-contiguous tracts of land in the lease sales. The possibility of strategic interactions thus poses a concern to policy-makers and affects the optimal government policy.

Since 1954, the U.S. government has leased tracts from its federal lands in the Gulf of Mexico to firms interested in offshore petroleum production by means of a succession of lease sales. A lease sale is initiated when the government announces that an area is available for exploration, and nominations are invited from firms as to which tracts should be offered for sale. In a typical lease sale, over a hundred tracts are sold simultaneously in separate first-price, sealed-bid auctions. Many more tracts are nominated than are sold, and the nomination process probably conveys little or no information (Porter, 1995). A tract is typically a block of 5000 acres or 5760 acres. The size of a tract is often less than the acreage required to ensure exclusive ownership of any deposits that may be present (Hendricks & Kovenock, 1989), and tracts within the same area may be located over a common pool (Hendricks & Porter, 1993). To date, the largest petroleum field spanned 23 tracts. Depending on water depth, 57-67 percent of the fields spanned more than one tract and 70-79 percent spanned three or fewer tracts (Marshall Rose, Minerals Management Service, personal communication, 31 March 2005). Because neighboring tracts of land may share a common pool of petroleum reserve, information and extraction externalities that lead firms to interact strategically may be present. As a consequence, petroleum production on the federal leases may be inefficient.

In this paper, I analyze whether a firm's investment timing decisions and profits depend on the decisions of firms owning neighboring tracts of land. Do the positive information externalities and negative extraction externalities have any net strategic effect?

To answer this question, I develop and estimate a structural econometric model of the firms' multi-stage investment timing game, and apply the model to data on petroleum production on U.S. federal lands in the Gulf of Mexico.

The research presented in this paper is important for several reasons. First, an empirical analysis of investment timing decisions enables one to examine whether the strategic interactions that are predicted in theory actually occur in practice. Second, the estimation of strategic interactions, especially those that arise in dynamic decision-making, is of methodological interest. Third, my results have implications for leasing policy: if the strategic effects and externalities turn out to be large, then the program by which the U.S. government leases tracts to firms may be inefficient, and possible modifications should be considered.[4]

The results from the structural econometric model do not indicate that externalities from exploration have any net strategic effect. A firm's profits from development do not depend significantly on the exploration decisions of its neighbor. In contrast, externalities from development do have a net strategic effect. A firm's real profits

[4] Even in the absence of externalities, the U.S. federal leasing program may still be inefficient because the government collects royalties without sharing costs. However, externalities would lead to additional inefficiencies.

from developing increase when its neighbor develops, perhaps because this is a signal to the firm that the neighbor's exploratory efforts were successful, and therefore that there may be deposits present.

There are several possible explanations why the results reject strategic, non-cooperative behavior during exploration. One is that the tract sizes are large enough that cross-tract externalities are insignificant or even nonexistent. A second is that cross-tract externalities exist, but firms owning neighboring tracts cooperate to jointly internalize the inefficient externalities they impose on each other, for example by forming joint ventures in exploration or by consolidating their production rights through purchase or unitization. A third is that cross-tract externalities are significant, but the positive information externality exactly cancels the negative extraction externality, resulting in zero net strategic effects. To distinguish among these three explanations for the lack of strategic interactions during exploration, I estimate the strategic interactions by tract size.

According to the estimation by tract size, the importance of strategic interactions depends on tract size. As expected, strategic interactions are more likely to take place on smaller tracts, where the externalities are more acute. For small tracts, the effect of having a neighboring tract explored reduces real profits by about 26 million dollars, while having a neighboring tract developed raises real profits by about 3.5 million dollars. On the small tracts, the externalities cost about 17 million real dollars per tract developed. When the tract size is large enough, the net strategic effects of the externalities from both exploration and development disappear. Also as predicted by theory, during exploration the relative importance of the extraction externality with respect to the information externality is greater on small tracts than on large tracts; on large tracts, the two externalities cancel each other out.

The results suggest that, by selling predominantly large tracts, the federal government has avoided additional inefficiencies in petroleum production that may have resulted from non-cooperative strategic interactions.

The balance of the paper proceeds as follows. I begin with a brief review of the relevant literature in Section 2. Section 3 presents the data and model and Section 4 explains the econometric estimation. Section 5 analyzes the results. Section 6 concludes.

2 Contributions to Existing Literature

The exploration timing game in offshore petroleum production in the Gulf of Mexico has been examined in a seminal series of papers by Kenneth Hendricks, Robert Porter and their co-authors (see e.g. Hendricks & Kovenock, 1989; Hendricks & Porter, 1993; Hendricks & Porter, 1996). These papers focus on the information externality associated with exploratory drilling. They analyze this externality and the learning and strategic delay that it causes by developing theoretical models of the exploration timing game. In addition, Hendricks and Porter (1993, 1996) calculate the empirical drilling hazard functions for cohorts in specific areas, and study the determinants of the exploration timing decision and of drilling outcomes. According to their results, equilibrium predictions of plausible non-cooperative models are reasonably accurate and more descriptive than those of cooperative models of drilling timing. Building upon their work, Lin (2009) further analyzes strategic interactions during exploration by estimating a reduced-form discrete response model of a firm's exploration timing decision that uses variables based on the timing of a neighbor's lease term as instruments for the neighbor's decision, and finds that there are no net strategic interactions during exploration.

The structural econometric analysis presented in this paper improves upon the existing literature on the exploration timing game in offshore petroleum production in several ways. First, unlike the theoretical models and reduced-form empirical analyses conducted by Hendricks, Porter and their co-authors, a structural approach yields estimates of the structural parameters of the discrete choice dynamic game. With these structural parameters, one can identify the effects of a neighbor's exploration and development decisions on the profits a firm would get from developing its tract.

A second way in which this paper contributes to the existing literature on the information externality in offshore petroleum production is that it combines the externality problem with real options theory. Oil production is a multi-stage process involving sequential investment decisions. Since the decision to explore a reserve entails an irreversible investment, the value of an unexplored reserve is the value of the option to invest in exploration. Similarly, the value of an explored but undeveloped reserve is the value of the option to invest in development. There is thus an option value to waiting before making either investment because the value of a developed reserve can change, either because exogenous conditions such as the oil price might change, or because there is a chance that neighboring firms might explore or develop first. Moreover, because these two types of investment are made sequentially, they act as compound options: completing one stage gives the firm an option to complete the next (Dixit & Pindyck, 1994).

While literature on the financial theory of option valuation is abundant, structural models applying the theory to the oil production process that account for strategic considerations have yet to be developed. Hurn and Wright (1994) test reduced-form implications of the theory via a hazard model, but neither estimate a structural model nor account for possible strategic interactions. Paddock, Siegel and Smith (1998) compare the option valuation estimate of the market value of selected offshore petroleum tracts with estimates from other valuation methods and with the winning bids, but do not account for either information externalities or extraction externalities. Similarly, Pesaran (1990) estimates an intertemporal econometric model for the joint determination of extraction and exploration decisions of a "representative" profit-maximizing oil producer, but does not examine the case of multiple producers that may interact strategically. Kellogg (2010) analyzes the extent to which firms delay irreversible investments following an increase in the uncertainty of their environment using detailed data on oil well drilling in Texas and expectations of future oil price volatility derived from the NYMEX futures options market, and finds that oil companies respond to changes in expected price volatility by adjusting their drilling activity by a magnitude consistent with the optimal response prescribed by theory.

The third innovation this paper makes to the existing literature on the information externality in offshore petroleum production is that while the existing literature focuses exclusively on externalities that arise during exploratory drilling, the model in this paper allows for extraction externalities as well as information externalities that arise during both exploration and development. If firms do indeed learn about the value of their own tracts from the actions of their neighbors, then one would expect firms to update their own beliefs not only if their neighbors begin exploratory drilling, but also if, after having already begun exploring, the neighbors then decide to install a production platform. That a neighbor has decided to begin extracting after it explored should be at least as informative as the initiation of exploration in the first place. Furthermore, extraction externalities are another form of spillover that is not accounted for by previous studies of the investment timing game, and, unlike the information externality, is one that may have a negative effect on a firm's profits.

In addition to the literature on the information externality, a second branch of related literature is that on econometric models of discrete dynamic games (see e.g. Aguirregabiria & Mira, 2007; Bajari, Benkard & Levin, 2007 and references therein). In particular, this paper applies a method developed by Pakes, Ostrovsky and Berry (2007) for estimating parameters of discrete dynamic games such as those involving firm entry and exit. This paper builds upon the work of Pakes et al. (2007) in several ways. First, unlike their paper, which uses simulated data, this paper estimates a discrete dynamic game using actual data. Second, while the entry and exit decisions they examine are two independent investments, the exploration and development decisions I examine are sequential investments: the decision to invest in development can only be made after exploration has already taken place. Thus, unlike the one-stage entry and exit games, the investment timing game is a two-stage game. The sequential nature of the investments is an added complexity that I address in my econometric model. Third, whereas the estimators Pakes et al. propose are for infinite-horizon dynamic games, the exploration stage of petroleum production is a finite-horizon dynamic optimization problem: firms must begin exploration before the end of the five-year lease term, or else relinquish their lease. As a consequence, an appropriate modification to the estimation algorithm is required. A fourth innovation I make is that, unlike Pakes et al., I estimate parameters in the profit function in addition to the parameters in the distribution of the private information. The task of estimating these additional parameters requires the use of additional moment conditions.

There are several advantages to using a structural model instead of the reduced-form approaches used by Hendricks and Porter (1993, 1996) and by Lin (2009). First, a structural model enables the estimation of all the structural parameters of the underlying dynamic game. These parameters include not only those governing the relationship between various state variables and the profits of firms, but also parameters governing the distribution of tract-specific private information.

A second advantage of a structural model is that it addresses the endogeneity problems without the need for instruments. Measuring neighbors' effects is difficult owing to two sources of endogeneity. One source is the simultaneity of the strategic interaction: if tract i is affected by its neighbor j, then tract j is affected by its neighbor i. The other arises from spatially correlated unobservable variables (Manski, 1993; Manski, 1995; Robalino & Pfaff, 2005). Because the structural model is based on the equilibrium of the underlying dynamic game, however, it addresses the simultaneity problem directly by explicitly modeling the firms' strategies. Moreover, the problem of spatially correlated unobservables can be addressed by interpreting the profits in the model as expected profits conditional on observables, where the expectation is taken over the correlated unobservables. In this case, the coefficients on the strategic variables measure the expected effect of the strategic variables, where the expectation is taken over the correlated unobservables. Thus, the model is still able to separately identify the (expected) strategic interaction from the correlated unobservable. As shown by results reported in the online Appendix of Monte Carlo experiments analyzing the effect of a state variable that is observed by the firms when they make their decisions but unobservable to the econometrician (i.e., a common shock), the bias introduced by spatially correlated unobservables is small. This is consistent with Pakes et al. (2007), who find that the bias from serially correlated common shocks is small.

A third advantage to a structural model is that it enables one to estimate how a firm's profits are affected by the decisions of its neighbors; the sign of the effect indicates the net sign of the information and extraction externalities. One is therefore able to quantify the net effects of the externalities.

A fourth advantage is that the structural model enables one to explicitly model each of the stages of the multi-stage dynamic decision-making problem faced by petroleum-producing firms. As a consequence, the analysis of strategic interactions in this paper is more complete than that of previous papers because it incorporates the second stage of petroleum production – development – along with the first stage – exploration – not only by allowing for strategic interactions in both stages, but also by linking the decisions made in each stage together in one integrated, multi-stage model that recognizes that decisions made in the first stage depend on the value of advancing to the second stage.

A fifth advantage is that with a structural model one can simulate what would happen if the tracts in a neighborhood were owned and managed by a single firm that internalizes the externalities.

3 A Model of the Investment Timing Game

I use a data set on federal lease sales in the Gulf of Mexico between 1954 and 1990 compiled by Kenneth Hendricks and Robert Porter from U.S. Department of

Interior data. There are three types of tracts that can be offered in an oil and gas lease sale: wildcat, drainage, and developmental. Wildcat tracts are located in regions where no exploratory drilling has occurred previously and therefore where the geology is not well known. Exploration on wildcat tracts entails searching for a new deposit. In contrast, both drainage and developmental tracts are adjacent to tracts on which deposits have already been discovered; developmental tracts, in addition, are tracts that have been previously offered in an earlier lease sale but either whose previous bids were rejected as inadequate or whose leases were relinquished because no exploratory drilling was done (Porter, 1995).

Time t denotes the number of years after the lease sale date. The time of exploration is the year of the tract's first spud date. The time of development is the year of the tract's first production date.[5] It is possible that development begins in the same year as exploration does. Firms must begin exploration before time T, the length of the lease term, or else relinquish their lease. Let the "lease term time" τ_{kt} of market k at time t be given by:

$$\tau_{kt} = \begin{cases} t & \text{if} \quad t = 0, 1, ..., T-1 \\ T & \text{if} \quad t \geq T \end{cases} .$$

I focus my attention on wildcat tracts offshore of Louisiana and Texas that were auctioned between 1954 and 1979, inclusive. I do so for several reasons. First, my restrictions are similar to those made by Hendricks and Porter (1996), thus enabling me to best compare my results with theirs. Second, since wildcat tracts are tracts on which no exploratory drilling has occurred previously, information externalities are likely to be most acute. Third, because the data set only contains production data up until 1990, the restriction to tracts sold before 1980 eliminates any censoring of either drilling or production.[6] Additional restrictions I impose for a tract to be included in my data set are that it must be a tract for which location data is available, for which the first exploration occurred neither before the sale date nor after the lease term,[7]

[5]More specifically, for both the time of exploration and the time of development, I take the floor of the number of years since the lease began.

[6]Another reason to focus on the earlier lease sales is that post-auction lease transfers occurred less frequently in the past (Porter, 1995; John Rodi, Minerals Management Service, personal communication, 8 May 2003; Robert Porter, personal communication, 21 May 2003).

[7]It is possible for a lease to receive a suspension of production (SOP) or suspension of operations (SOO) which will extend the life of the lease beyond its primary term (Jane Johnson, Minerals Man-

and for which production did not occur before exploration.

In my model of the investment timing game, each "market" k consists of an isolated neighborhood of adjacent tracts i that were each leased to a petroleum-producing firm on the same date. From the wildcat tracts in my data set, I select tracts that fall into two-tract markets. In order for two tracts to qualify as a market, the two tracts must be within 6 miles north and south of each other or 6 miles east and west of each other,[8] both the sale dates and the lease terms must be the same for both tracts, the tracts must be owned by different firms, and no other wildcat tracts from the same sale date can be within 6 miles of either tract. There are 87 such markets in my data set. Figure 1 maps out the tracts used.

I restrict the size of the market to two tracts for two reasons. First, limiting the market size to two minimizes the state space. This is because the number of possible combinations of state variables, which is the product of the cardinality of the supports of each of the state variables, is quadratic in the market size. Second, when there are only two tracts in the market, each tract is equidistant to all its neighbors. Since the federal tracts in the Gulf of Mexico form a grid, the next sensible size of a market is four. With four tracts, however, diagonal tracts are not as close together as tracts that share a side are; as a consequence, the distance between each pair of neighbors in the market is not the same. It is plausible that firms may weight the behavior of their neighbors by their distance; when neighbors are no longer symmetric, the econometric model becomes more complicated. To ensure that the two tracts in the market are adjacent to each other, I require the two tracts to be within 6 miles north and south of each other or 6 miles east and west of each other because the maximum tract size is 5760 acres, or 3 miles by 3 miles, and therefore the tracts have to be within 6 miles of each other to be adjacent. According to Table 1, which compares summary statistics for tracts in the data set that fall in the two-tract markets with tracts in the data set that do not, the tracts in the two-tract markets appear representative of all the wildcat tracts in the data set. The number of years to exploration, the number of years to development, revenue, and gross profits for tracts in the two-tract markets are similar to those not in the two-tract markets.

---

agement Service, personal communication, October 29, 2003). Exploratory drilling first occurred after the lease term on 77 (or 3.1 %) of the 2481 wildcat tracts sold before 1980.

[8] I convert latitude and longitude to miles using the following following factors from the Lousiana Sea Grant web site (http://lamer.lsu.edu/classroom/deadzone/changedistance.htm): 1 minute longitude in Louisiana offshore = 60.5 miles; 1 minute latitude = 69.1 miles.

The panel spans the years 1954 to 1990. A market enters the panel when its tracts are sold. If both tracts are eventually developed, the market exits the panel when the second tract to develop first begins development. If neither tract has explored by the end of the five-year lease term, the market exits the panel when the leases expire.

For each market k, the state of the market t years after the leases began is given by a vector Ω_{kt} of discrete and finite-valued state variables that are observed by all the firms in market k and as well as by the econometrician. Let θ denote the vector of parameters to be estimated.

At the beginning of each period t, the owner of each tract i must make one of two investment decisions. If tract i has not been explored before time t, its owner must decide whether to invest in exploration at time t. If tract i has been explored but has not been developed before time t, its owner must decide whether to invest in development at time t. For each period t, all firms make their time-t investment decisions simultaneously.

Each firm's time-t investment timing decision depends in part on the state of the market $\Omega_{kt} \equiv (N_{kt}, X_{kt}, \tau_{kt})$, which can be decomposed into endogenous state variables N_{kt}, exogenous profit-shifting state variables X_{kt}, and the lease term time τ_{kt}. Investment decisions depend on N_{kt} and X_{kt} because these state variables are assumed to affect profits. Because of the finite-horizon nature of the firm's exploration investment problem, the finite-valued and exogenous lease term time τ_{kt} affects investment decisions as well, as will be explained below. In the present model, there are two endogenous state variables N_{kt}: the total number of tracts in market k that have been explored before time t, and the total number of tracts in market k that have been developed before time t. These endogenous state variables capture the strategic component of the firms' investment timing decisions. The exogenous state variables X_{kt} include the drilling cost and the oil price and are assumed to evolve as a finite state first-order Markov process: $X_{k,t+1} \overset{iid}{\sim} F_X(\cdot|X_{kt})$. In other words, the next period's value $X_{k,t+1}$ of the exogenous state variables are assumed to be independently and identically distributed (i.i.d.) with a probability distribution that depends only on the time-t realization X_{kt} of the exogenous state variables, and not additionally on what happened before time t (Dixit & Pindyck, 1994).[9]

[9] The lease term time τ_{kt} evolves as a finite state first-order Markov process as well. I include this exogenous finite-valued variable τ_{kt} as a separate argument distinct from X_{kt} both because it

I use three exogenous state variables X_{kt}; these variables were chosen based on considerations of state space and data availability. The first exogenous state variable is the discretized average winning bid per acre over the two tracts in market k at time t. Because the winning bid is a measure of the value of the tract, the average winning bid over the tracts in the market captures any fixed market-specific variables such as geological structures that may affect profits. To construct this variable I average the winning bids per acre over the two tracts in the market, and then discretize the average into three bins: 0 = low (0 to 1 thousand 1982 \$/acre), 1 = medium (1 thousand to 5 thousand 1982 \$/acre), and 2 = high (over 5 thousand 1982 \$/acre). One expects that profits would increase in the value of the tract, and therefore that the coefficient on the winning bid in the development profit function is positive.

The second exogenous state variable is the discretized real drilling cost at time t. I use data on annual drilling costs from the American Petroleum Institute's *Joint Association Survey of the U.S. Oil & Gas Producing Industry* for the 1969-1975 data and its *Joint Association Survey on Drilling Costs* for the 1976-1990 data. The cost is average cost per well over all offshore wells (oil wells, gas wells, dry holes), in nominal dollars. I convert the nominal costs to real costs in 1982-1984 dollars using the consumer price index (CPI). I discretize the real drilling cost into two bins: 0 = low (0 to 2.5 million 1982-1984 \$/well), 1 = high (over 2.5 million 1982-1984 \$/well). The drilling cost is a measure of both exploration costs and development costs, and therefore enters into both the extraction profit function and the development profit function. The discretized drilling cost is incremented by one so that costs are non-zero even when they fall in the low bin. Figure 2 plots the real drilling cost data, along with the bins.[10] The expected sign of the coefficient on costs in the exploration profit function is negative; higher costs should lower exploration profits. Similarly, the expected sign of the coefficient on costs in the development profit function is negative as well.

The third exogenous state variable is the discretized real oil price. I use the U.S. average crude oil domestic first purchase price from the EIA Annual Energy Review and deflate the time series to 1982-1984 dollars per barrel using the CPI. I discretize

does not affect profits and also to elucidate my later exposition of the finite-horizon nature of the exploration stage.

[10] Before 1969, the real drilling cost is assumed to fall into the low bin. The 1982 real drilling cost, which was unavailable because the 1982 issue of the *Joint Association Survey on Drilling Costs* was out of print, is assumed to fall in the high bin.

the real oil price into three bins: 0 = low (0 to 13 1982-1984 \$/barrel), 1 = medium (13 to 25 1982-1984 \$/barrel), and 2 = high (over 25 1982-1984 \$/barrel). Figure 3 plots the real oil price, along with the bins. The expected sign on oil price in the development profit function is positive: higher oil prices should increase revenues.

There are two endogenous state variables N_{kt}: the total number of tracts in the market that have been explored and the total number of tracts in the market that have been developed. Because there are two tracts in each market, there are three possible values for the the total number of tracts that have been explored: 0, 1 and 2. As for the total number of tracts that have been developed, because a market ends once both tracts have been developed, there are only two possible values: 0 and 1.

Table 2 presents summary statistics for the panel data used for the structural estimation. There are 1041 observations spanning 174 tracts and 87 markets. The markets range in duration from 2 years to 36 years, with an average length of 17.92 years (s.d. = 11.17). Of the 174 tracts, 122 were eventually explored and 66 were eventually developed. The average number of years to exploration, conditional on exploring, is 1.21 (s.d. = 1.42). The average number of years to development, conditional on developing, is 5.79 (s.d. = 3.51). For the 66 tracts that developed, the predicted ex post revenues, as calculated by Hendricks, Porter and Boudreau (1987), range from \$22,000 to \$298 million, with an average of \$49.34 million (s.d. = 65.53 million). The real gross profits from development, which are the predicted ex post revenues times the government royalty rate minus costs, but not net of the bid, also as calculated by Hendricks, Porter and Boudreau (1987), range from -\$38.10 million to \$18.80 million, with an average of -\$10.83 million (s.d. = 9.57 million). Table 2 also provides the summary statistics for the state variables in the panel. The number of possible combinations of state variables is the product of the cardinality of the supports of each of the state variables, or $3 \times 2 \times 3 \times 2 \times 3 = 108$.[11]

The maximum tract size, as stipulated by a provision in section 8(b) of the Outer Continental Shelf Lands Act (OCSLA), 43 U.S.C. 1337(b)(1), is 5760 acres, or 3 miles by 3 miles. The distribution of tract sizes is presented in Figure 4. Most tracts are

[11] Both the small-sample bias and variance of the estimator depend on the variances of the non-parametric component. For a given sized data set, the larger the number of states, the fewer the number of observations per state and the larger the variance in the first-stage estimates is likely to be (Pakes et al, 2007). Because the sample is small, it is assumed that lots of transitions cannot happen because they are never observed in the data. Even with the small sample, however, Monte Carlo results reported in the online Appendix show that the estimators recover the actual parameter values fairly well.

either 2500 acres, 5000 acres or 5760 acres in size. The mean tract size is 4460 acres (s.d. = 1300), and the median tract size is 5000 acres. Table 3 presents summary statistics by tract size. Large tracts are defined as tracts that are greater than or equal to 5000 acres in size; small tracts are defined as tracts that are less than 5000, 4000 or 3000 acres in size, respectively.[12] The distributions of the variables appear to be similar across tract sizes. Small tracts make up a small percentage of the tracts sold: of the tracts used in the structural estimation, only 37%, 27%, and 24% of tracts were less than 5000 acres, 4000 acres and 3000 acres in size, respectively.

In addition to the publicly observable state variables Ω_{kt}, each firm's time-t investment timing decision also depends on two types of shocks that are private information to the firm and unobserved by either other firms or by the econometrician. The first source of private information is a pre-exploration shock μ_{ikt} to an unexplored tract i at time t. This pre-exploration shock, which is only observed by the firm owning tract i, represents any and all private information that affects the exploration investment decision made on tract i at time t. Such private information may include, for example, idiosyncratic shocks to exploration costs and the outcome of the post-sale, pre-exploration seismic study conducted on tract i at time $t-1$.[13] Following Pakes, Ostrovsky and Berry (2007), assume that the pre-exploration shock μ_{ikt} is an independently and identically distributed random variable with an exponential distribution and mean σ_μ. That is, $\mu_{ikt} \overset{iid}{\sim} exponential(\sigma_\mu)$.

The second source of private information is a pre-development shock ε_{ikt} to an explored but undeveloped tract i at time t. This pre-development shock, which is only observed by the firm owning tract i, represents any and all private information that affects the development investment decision made on tract i at time t. Such private information may include, for example, idiosyncratic shocks to development costs and the outcome of the exploratory drilling conducted on tract i at time $t-1$. Following Pakes, Ostrovsky and Berry (2007), assume that the pre-development shock ε_{ikt} is

[12] The sum of the number of markets with tracts greater than or equal to 5000 acres and the number of markets with tracts less than 5000 acres is three fewer than the total number of markets because these three markets each had one tract that was greater than or equal to 5000 acres and another tract that was less than 5000 acres.

[13] Firms conduct and analyze seismic studies in order to help them decide whether or not to begin exploratory drilling (John Shaw, personal communication, 18 April 2003; Bob Dye, Apache, personal communication, 21 January 2004; Jon Jeppesen, Apache, personal communication, 21 January 2004; Mark Bauer, Apache, personal communication, 21 January 2004; Billy Ebarb, Apache, personal communication, 22 January 2004).

an independently and identically distributed random variable with an exponential distribution and mean σ_ε. That is, $\varepsilon_{ikt} \overset{iid}{\sim} exponential(\sigma_\varepsilon)$. In addition, assume that the pre-exploration shocks μ_{ikt} and the pre-development shocks ε_{ikt} are independent of each other.[14] All distributions are common knowledge.

In the absence of strategic considerations, the firm owning tract i would base its investment timing decisions on only the exogenous state variables X_{kt}, the lease term time τ_{kt}, and the private shocks μ_{ikt} and ε_{ikt}. To derive its dynamically optimal investment policy, it would solve a single-agent dynamic programming problem.

If information and extraction externalities were present, however, then strategic considerations would become important. As a consequence, the exploration and development investment decisions of the firm owning tract i in market k would depend on the exploration and development investment decisions of the firms owning the other tracts in market k. In other words, the firm owning tract i would base its investment timing decisions not only on the exogenous state variables X_{kt}, the lease term time τ_{kt}, and the private shocks μ_{ikt} and ε_{ikt}, but also on the endogenous state variables N_{kt} as well, namely the total number of tracts in its market k that have been explored before time t and the total number of tracts in market k that have been developed before time t. Each firm would then no longer solve merely a single-agent dynamic programming problem, but rather a multi-agent dynamic game.

The equilibrium concept used in the model is that of a Markov perfect equilibrium. Each firm is assumed to play a Markov "state-space" strategy: the past influences current play only through its effect on the state variables. A firm's dynamically optimal investment policy is then the Markov strategy that it plays in the Markov perfect equilibrium, which is a profile of Markov strategies that yields a Nash equilibrium in every proper subgame (Fudenberg & Tirole, 1998).

[14] The assumptions that both types of shocks are i.i.d. and independent of each other, while restrictive, are needed in order for the estimation technique used in this paper to work. If either type of shock were serially correlated (or if, at the extreme, there were tract fixed effects), then firms would base their decisions not only on the current values of the state variables and of their shocks, but also on past values of the state variables and shocks as well. The state space would then be too large. If the distribution of the pre-development shock ε_{ikt} depended on the realization of the pre-exploration shock μ_{ikt} (e.g., the μ_{ikt} at the time of exploration), then μ_{ikt} would be a state variable in the development stage of production. As a consequence, the econometrician would need to observe μ_{ikt}, which she does not. The i.i.d. assumption is reasonable if the shocks are interpreted to encompass all idiosyncratic factors affecting investment decisions, including managerial shocks and technological shocks. Moreover, since one of my state variables is the average winning bid, which is a measure of tract value, it is reasonable to assume that, conditional on tract value, shocks are i.i.d.

While each firm's time-t investment decision depends on both the publicly available endogenous and exogenous state variables Ω_{kt} as well as the firm's own private information μ_{ikt} or ε_{ikt}, its perception of its neighbor's time-t investment decisions depend only on the publicly observable state variables Ω_{kt}. This is because, owing to the above assumptions on the observable state variables and on the unobservable shocks, firms can take expectations over their neighbors' private information.[15] In equilibrium, firms' perceptions of their neighbors' investment probabilities should be consistent with those that are actually realized (Starr & Ho, 1969).

The model has at least one Markov perfect equilibrium, and each equilibrium generates a finite state Markov chain in Ω_{kt} tuples (Pakes, Ostrovsky & Berry, 2007).[16] Although model assumptions do not guarantee a unique equilibrium, they do insure that there is only one set of equilibrium policies that is consistent with the data generating process. It is thus possible to use the data itself to pick out the equilibrium that is played. For large enough samples, the data will pick out the correct equilibrium and the estimators for the parameters in the model will be consistent (Pakes, Ostrovsky & Berry, 2007).[17]

The firm's dynamic decision-making problem is as follows. The first-stage problem is to determine the optimal policy for investment in exploration. Because firms must begin exploration before the end of their lease term, or else relinquish their lease, this is a finite-horizon problem. As a consequence, firms' decisions will depend not only on the profit-shifting state variables N_{kt} and X_{kt}, but also on time t. However, since firms can only make exploration decisions at the beginning of periods $t = 0, ..., T-1$, the time dependence only applies until time $t = T-1$, after which exploration can no longer begin and the endogenous variable that counts the total number of tracts in the market that have been explored stays constant. It is for this reason that the exogenous and finite-valued state variable "lease term time" τ_{kt} captures the entire time dependence of the problem.

The second stage of the firm's dynamic decision-making problem is to determine the optimal timing for investing in the development of a tract that has already been

[15] While each firm plays a pure strategy, from the point of view of their neighbors, they appear to play mixed strategies. Thus, as with Harsanyi's (1973) purification theorem, a mixed distribution over actions is the result of unobserved payoff perturbations that sometimes lead firms to have a strict preference for one action, and sometimes a strict preference for another.

[16] A Markov chain is a Markov process on a finite state space (Stokey, Lucas & Prescott, 1989).

[17] This assumes that the same equilibrium is played in each market. If a mixed strategy equilibrium is played, then it is assumed that the same mixed strategy equilibrium is played in each market.

explored. This second-stage problem has both a finite-horizon component and an infinite-horizon component. A firm's development strategy depends in part on its perceptions of the future exploration policies of the firms in the market. Since exploration policies depend on time until time $t = T - 1$, this means that perceptions, and therefore development strategies, will depend on time for $t < T$. As a consequence, the dynamic programming problem for time $t < T$ is a finite-horizon problem. However, because the lease term only applies to the exploration stage of production, and because the endogenous variable that counts the total number of tracts in the market that have been explored – a variable that depends on the time-dependent exploration policies of the firms in the market – stays constant after the lease term expires, the dynamic programming problem for the development stage from time T onwards is an infinite-horizon problem that does not depend on time. Thus, once again, the lease term time τ_{kt} sufficiently captures the entire time dependence of the problem.

The firm's sequential investment problem is a two-stage optimization problem, and can be solved backwards using dynamic programming (Dixit & Pindyck, 1994). In the second, or development, stage of oil production, a firm with an explored but undeveloped tract i must decide if and when to invest in a production platform. Assume that the profit $\pi^d(\Omega_{kt}, \varepsilon_{ikt}; \theta)$ that a firm will get after developing tract i at time t can be separated into a deterministic component and a stochastic component as follows:

$$\pi^d(\Omega_{kt}, \varepsilon_{ikt}; \theta) = \pi_0^d(\Omega_{kt}; \theta) + \varepsilon_{ikt} \ , \tag{1}$$

where the deterministic component of profit is linear in the publicly observable state variables:

$$\pi_0^d(\Omega_{kt}; \theta) \equiv N_{kt}' \gamma_N + X_{kt}' \gamma_X \ , \tag{2}$$

and where the stochastic component is the privately observed pre-development shock ε_{ikt}. The development profit is therefore independent of time (and lease term time) except through the state variables (N_{kt}, X_{kt}) and the shock ε_{ikt}. This development profit $\pi^d(\Omega_{kt}, \varepsilon_{ikt}; \theta)$ is the present discounted value of the entire stream of profits from developing a tract.

If there were additional market state variables that affected profits but were unobserved by the econometrician, then $\pi^d(\Omega_{kt}, \varepsilon_{ikt}; \theta)$ can be interpreted as the expected profits conditional on the available information Ω_{kt} (Pakes et al., 2007). Under this interpretation, spatially correlated unobservables do not pose a concern. If these unobservable state variables were observed by the firms but not by the econometrician, Monte Carlo experiments presented in Section 4 show that the bias they may introduce is small.

The development profit function is a function of variables that affect the development profit, including oil price and variables that are indicators of the amount of oil that can be extracted. Variables that are indicators of the amount of oil that can be extracted include the average winning bid per acre as well as the number of tracts in the market that have explored and the number of tracts in the market that have developed. The number of tracts neighboring tract i that have explored and the number of tracts neighboring tract i that have developed are potential indicators of how much oil is in the ground under tract i because they indicate whether or not a neighbor thought it was worthwhile to explore and/or develop its tract, which might share common geological features with firm i's tract (the information externality), and are also potential indicators of how much oil can be extracted from tract i because they indicate whether or not firm i will be competing for oil with its neighbor (the extraction externality).

Let $\gamma \equiv (\gamma_N, \gamma_X)$ denote the vector of all the coefficients in the development profit function. The coefficients γ_N in the profit function on the endogenous state variables N_{kt} – the total number of tracts in the market that have been explored and the total number of tracts in the market that have been developed – indicate whether and how one firm's profits depend on the production decisions of its neighbors. If a neighbor explores, then the state variable counting the total number of tracts in the market that have been explored increases by one and the value of the development profits increase by the value of its coefficient. Similarly, if a neighbor develops, then the state variable counting the total number of tracts in the market that have been developed increases by one and the value of the development profits increase by the value of its coefficient. The coefficients γ_N on the endogenous variables thus measure the net effects of the information and extraction externalities, and therefore indicate whether firms interact strategically on net. Positive values of the coefficients γ_N would indicate that the information and extraction externalities were positive on net, and therefore that the information externality was dominant. Negative values would indicate that the externalities were negative on net, and therefore that the extraction externality was dominant. Moreover, one would expect each coefficient γ_N to be less than the cost of development. Otherwise, if the coefficients were greater than or equal to the development cost, this would mean that having a neighbor explore or develop would offset the cost of development, making development essentially costless. The effects of strategic interaction are unlikely to be that large.

The value V^e of an explored but undeveloped tract i in market k at time t is given by:

$$V^e(\Omega_{kt}, \varepsilon_{ikt}; \theta) = \max\{\pi^d(\Omega_{kt}, \varepsilon_{ikt}; \theta), \beta V^{ce}(\Omega_{kt}; \theta)\}, \tag{3}$$

where $\beta \in (0, 1)$ is the discount factor and $V^{ce}(\Omega_{kt}; \theta)$ is the continuation value to waiting instead of developing at time t. For the structural estimation, I set the discount factor β to 0.9. The continuation value to waiting is the expectation over the state variables and shocks of next period's value function, conditional on not developing this period:

$$V^{ce}(\Omega_{kt}; \theta) = E\left[V^e(\Omega_{k,t+1}, \varepsilon_{i,k,t+1}; \theta)|\Omega_{kt}, I^d_{ikt} = 0\right], \tag{4}$$

where I^d_{ikt} is an indicator for whether development began on tract i at time t.

Let $g^d(\Omega_{kt}; \theta)$ denote the probability of developing an explored but undeveloped tract i at time t conditional on the publicly available information Ω_{kt} on time t, but not on the private information ε_{ikt}. The development probability $g^d(\Omega_{kt}; \theta)$ function represents a firm's perceptions of the probability that a neighbor owning an explored but undeveloped tract will decide to develop its tract in period t, given that the state of their market at time t is Ω_{kt}. Moreover, a firm's expectation of its own probability of development in the next period is simply the expected value of the next period's development probability, conditional on this period's state variables: $E[g^d(\Omega_{k,t+1}; \theta)|\Omega_{kt}]$.

Using the exponential distribution for ε_{ikt} and equation (1) for development profits, the continuation value $V^{ce}(\cdot)$ can be reduced to:

$$V^{ce}(\Omega_{kt}; \theta) = E[\beta V^{ce}(\Omega_{k,t+1}; \theta) + \sigma_\varepsilon g^d(\Omega_{k,t+1}; \theta)|\Omega_{kt}, I^d_{ikt} = 0], \tag{5}$$

and the development probability $g^d(\cdot)$ can be reduced to the following function of the continuation value, the state variables and the parameters:

$$g^d(\Omega_{kt};\theta) \;= \exp\left(-\frac{\beta V^{ce}(\Omega_{kt};\theta) - \pi_0^d(\Omega_{kt};\theta)}{\sigma_\varepsilon}\right). \tag{6}$$

In the first, or exploration, stage of oil production, a firm with an unexplored tract i must decide if and when to invest in exploratory drilling. Owing to the sequential nature of the investments, the publicly observable deterministic component of the payoff $\pi_0^e(\cdot)$ to exploring in the first stage is equal to the expected value of having an explored but undeveloped tract in the second stage, net the cost of exploration $c^e(\cdot)$:

$$\pi_0^e(\Omega_{kt};\theta) \equiv E_\varepsilon\left[V^e(\Omega_{kt},\varepsilon_{ikt};\theta)|\Omega_{kt}\right] - c^e(\Omega_{kt};\theta), \tag{7}$$

where the exploration cost is assumed to be linear in the exogenous cost-shifting state variables:[18]

$$c^e(\Omega_{kt};\theta) = -X'_{kt}\alpha \;. \tag{8}$$

In particular, the exploration cost is assumed to be the following function of the discretized drilling cost $drill\_cost_t$:

$$c^e(\Omega_{kt};\theta) = -\alpha\cdot(drill\_cost_t + 1)\,, \tag{9}$$

where α is now a scalar, so that exploration profits are:

$$\pi^e(\Omega_{kt},\mu_{ikt};\theta) = E_\varepsilon\left[V^e(\Omega_{kt},\varepsilon_{ikt};\theta)|\Omega_{kt}\right] + \alpha\cdot(drill\_cost_t + 1) + \mu_{ikt}. \tag{10}$$

Assume that the actual payoff $\pi^e(\cdot)$ to exploring tract i at time t also includes a privately observed stochastic component as well:

$$\pi^e(\Omega_{kt},\mu_{ikt};\theta) = \pi_0^e(\Omega_{kt};\theta) + \mu_{ikt}\;, \tag{11}$$

where the stochastic component is the pre-exploration shock μ_{ikt}.

The value V^n of an unexplored tract i in market k and time t is given by:

$$V^n(\Omega_{kt},\mu_{ikt};\theta) = \max\{\pi^e(\Omega_{kt},\mu_{ikt};\theta), \beta V^{cn}(\Omega_{kt};\theta)\}, \tag{12}$$

[18] I define costs with a negative sign so that the coefficients can be interpreted as coefficients in the exploration profit function. Variables that increase cost will decrease profit, and vice versa.

where $V^{cn}(\Omega_{kt};\theta)$ is the continuation value to waiting instead of exploring at time t. The continuation value to waiting is the expectation over the state variables and shocks of next period's value function, conditional on not exploring this period:

$$V^{cn}(\Omega_{kt};\theta) = E\left[V^{n}(\Omega_{k,t+1},\mu_{i,k,t+1};\theta)|\Omega_{kt},I^{e}_{ikt}=0\right], \tag{13}$$

where I^{e}_{ikt} is an indicator for whether exploration began on tract i at time t. The lease term imposes the following boundary condition:

$$V^{n}\left((N,X,T),\mu;\theta\right) = 0 \quad \forall N,X,\mu\ . \tag{14}$$

Let $g^{e}(\Omega_{kt};\theta)$ denote the probability of exploring an unexplored tract i at time t conditional on the publicly available information Ω_{kt} on time t, but not on the private information μ_{ikt}. As with the development probability, the current value of the exploration probability represents a firm's perceptions of the probability that a neighbor owning an unexplored tract will decide to explore its tract in period t, given that the state of their market at time t is Ω_{kt}; its expected value at time $t+1$ represents a firm's expectation of its own probability of exploration in the next period.

Using the exponential distribution for μ_{ikt} and equation (11) for exploration profits, the continuation value $V^{cn}(\cdot)$ to waiting instead of exploring can be reduced to:

$$V^{cn}(\Omega_{kt};\theta) = E\left[\beta V^{cn}(\Omega_{k,t+1};\theta)+\sigma_{\mu}g^{e}(\Omega_{k,t+1};\theta)|\Omega_{kt},I^{e}_{ikt}=0\right], \tag{15}$$

and the exploration policy function $g^{e}(\cdot)$ can be reduced to the following function of the continuation values, state variables and parameters:

$$g^{e}(\Omega_{kt};\theta) = \exp\left(-\frac{\beta V^{cn}(\Omega_{kt};\theta)-\left(\beta V^{ce}(\Omega_{kt};\theta)+\sigma_{\varepsilon}g^{d}(\Omega_{kt};\theta)\right)+c^{e}(\Omega_{kt};\theta)}{\sigma_{u}}\right). \tag{16}$$

Owing to the sequential nature of the investment decisions, the continuation value $V^{ce}(\cdot)$ and the investment probability $g^{d}(\cdot)$ from the development stage appear in the expression for the investment probability $g^{e}(\cdot)$ in the exploration stage.

The ex ante expected value of an unexplored tract at time $t=0$, where expectations are taken over the pre-exploration shock μ, is given by:

$$E_{\mu}\left[V^{n}(\Omega_{k0},\mu_{ik0};\theta)|\Omega_{k0}\right]=\beta V^{cn}(\Omega_{k0};\theta)+\sigma_{\mu}g^{e}(\Omega_{k0};\theta). \tag{17}$$

4 The Structural Econometric Model

The econometric estimation technique I use employs a two-step semi-parametric estimation procedure. It is an extension of the estimator proposed by Pakes, Ostrovsky and Berry (2007) to finite-horizon, multi-stage games. In the first step, the continuation values are estimated non-parametrically and these estimates are used to compute the predicted probabilities of exploration and development. In the second step, the parameters $\theta\equiv(\sigma_{\mu},\sigma_{\varepsilon},\gamma,\alpha)'$ are estimated by matching the predicted probabilities with the actual probabilities in the data. I will now describe each step in turn.

4.1 Step 1: Estimating continuation values and predicted probabilities

The first step entails computing the non-parametric[19] estimators $V^{ce}\widehat{(\Omega_{kt};\theta)}$ and $V^{cn}\widehat{(\Omega_{kt};\theta)}$ for the continuation values $V^{ce}(\Omega_{kt};\theta)$ and $V^{cn}(\Omega_{kt};\theta)$, respectively, given θ. To do so, historical empirical frequencies are used to estimate the elements of the Markov transition matrix governing the evolution of the finite-valued state variables from one period to the next. Estimators for the continuation values are subsequently derived from equations (5) and (15) using dynamic programming. These estimators are then substituted into equations (6) and (16) to obtain predicted probabilities for development and exploration, respectively.[20]

Formally, the non-parametric estimator $V^{ce}\widehat{(\Omega_{kt};\theta)}$ for $V^{ce}(\Omega_{kt};\theta)$ is derived from equation (5) and is computed as follows. For each period t, let each component of the

[19]The continuation values are non-parametric functions of the state variables Ω_{kt} conditional on the parameters θ.

[20]Rather than use historical empirical frequencies to estimate the Markov transition matrix, it is possible to compute an estimator for the matrix using the estimators for the exploration and development probabilities. However, because the latter, more complicated approach imposes a computational burden and because Pakes, Ostrovsky and Berry (2007) find that it did not improve the performance of their estimator, I choose the former, simpler approach.

vector $\overrightarrow{V_t^{ce}}$ be $V^{ce}(\Omega_{kt};\theta)$ evaluated at a different tuple of state variables. Similarly, for each period t, let each component of the vector $\overrightarrow{g_t^d}$ be $g^d(\Omega_{kt};\theta)$ evaluated at a different tuple of state variables. Finally, for each lease term time period τ, let M_τ^e be a transition matrix from the point of view of an owner of an explored but undeveloped tract who decides not to develop at time t. The element in the i^{th} row and j^{th} column is the probability that the state tuple next period will be the j^{th} tuple, given that the state tuple this period is the i^{th} tuple, given that the tract has already been explored but not yet developed at time t, and conditional on not developing at time t.

The estimator $V^{ce}\widehat{(\Omega_{kt};}\theta)$ is obtained from rewriting equation (5) in vector form:

$$\overrightarrow{V_t^{ce}} = M_\tau^e \left(\beta\overrightarrow{V_{t+1}^{ce}} + \sigma_\varepsilon \overrightarrow{g_{t+1}^d}\right). \tag{18}$$

The estimator is obtained after further substituting in the empirical average $\widehat{M_\tau^e}$ for M_τ^e. For $t \geq T$, since $\overrightarrow{V_t^{ce}} = \overrightarrow{V_{t+1}^{ce}}$ $\forall t \geq T$, we can solve for a fixed point $\widehat{V^{ce}}((N_{kt}, X_{kt}, T;\theta)$, which, from Blackwell's Theorem, is unique. To obtain the estimator of the value function for $t < T$, we then iterate backwards in time from $t = T$ using $\overrightarrow{V_T^{ce}} = \widehat{V^{ce}}((N_{kt}, X_{kt}, T;\theta)$ as a boundary condition. The predicted probability of development is then given by:[21]

$$g^d\widehat{(\Omega_{kt};}\theta) \;= \exp\left(-\frac{\beta V^{ce}\widehat{(\Omega_{kt};}\theta) - \pi_0^d(\Omega_{kt};\theta)}{\sigma_\varepsilon}\right). \tag{19}$$

The non-parametric estimator $V^{cn}\widehat{(\Omega_{kt};}\theta)$ for the continuation value to waiting instead of exploring is derived from equation (15) and is computed in a similar fashion as the estimator of $V^{ce}(\cdot)$. For each period t, let each component of the vector $\overrightarrow{V_t^{cn}}$ be $V^{cn}(\cdot)$ evaluated at a different tuple of state variables. For each period t, let each component of the vector $\overrightarrow{g_t^e}$ be $g^e(\cdot)$ evaluated at a different tuple of state variables. Finally, for each lease term time period τ, let M_τ^n be a transition matrix from the point of view of an owner of an unexplored tract who decides not to explore at time t. The element in the i^{th} row and j^{th} column is the probability that the state tuple next period will be the j^{th} tuple, given that the state tuple this period is the i^{th} tuple, given that the tract has yet to be explored, and conditional on not exploring at time t.

[21] In practice, I use: $g^d\widehat{(\Omega_{kt};}\theta) \;= \min\left\{\exp\left(-\frac{\beta V^{ce}\widehat{(\Omega_{kt};}\theta)-\pi_0^d(\Omega_{kt};\theta)}{\sigma_\varepsilon}\right), 1\right\}$.

The estimator $V^{cn}\widehat{(\Omega_{kt};\theta)}$ is obtained from rewriting equation (15) in vector form:

$$\overrightarrow{V_t^{cn}} = M_\tau^n \left(\beta \overrightarrow{V_{t+1}^{cn}} + \sigma_\mu \overrightarrow{g_{t+1}^e}\right). \tag{20}$$

Substituting in the empirical average $\widehat{M_\tau^n}$ for M_τ^n, we can solve backwards in time from the boundary condition $\overrightarrow{V_{T-1}^{cn}} = 0$ implied by equation (14) to obtain $\widehat{V^{cn}}(\Omega_{kt};\theta)$ for all $t \leq T-1$. The predicted probability of exploration is then given by:[22]

$$g^e\widehat{(\Omega_{kt};\theta)} = \exp\left(-\frac{\beta V^{cn}\widehat{(\Omega_{kt};\theta)} - \left(\beta V^{ce}\widehat{(\Omega_{kt};\theta)} + g^d\widehat{(\Omega_{kt};\theta)}\cdot\sigma_\varepsilon\right) + c^e(\Omega_{kt};\theta)}{\sigma_u}\right). \tag{21}$$

Owing to the sequential nature of the investment decisions, the estimators $V^{ce}\widehat{(\Omega_{kt};\theta)}$ and $g^d\widehat{(\Omega_{kt};\theta)}$ of the continuation value and the investment probability, respectively, from the development stage are needed to form the estimator of the investment probability $g^e\widehat{(\Omega_{kt};\theta)}$ in the exploration stage.

4.2 Step 2: Generalized Method of Moments

After obtaining estimates of the continuation values and predicted probabilities as functions of the state variables Ω_{kt} and the parameters θ in the first step, I estimate the parameters θ in the second step using generalized method of moments (GMM). The moments I construct involve matching the probabilities of exploration and development predicted by the model, as given by equations (21) and (19), with the respective empirical probabilities $\overline{g^e(\Omega_{kt})}$ and $\overline{g^d(\Omega_{kt})}$ in the data. I also form moments that match, for those tracts that developed, the expected development profits conditional on development predicted by the model with the actual average realized profits $\overline{\pi^d(\Omega_{kt})}$ in the data. Additional moments are constructed by interacting the above moments with the state variables. The moment function $\Psi(\Omega_{kt},\theta)$ is therefore:

[22] In practice, I use: $g^e\widehat{(\Omega_{kt};\theta)} = \min\left\{\exp\left(-\frac{\beta V^{cn}\widehat{(\Omega_{kt};\theta)} - \left(\beta V^{ce}\widehat{(\Omega_{kt};\theta)} + g^d\widehat{(\Omega_{kt};\theta)}\cdot\sigma_\varepsilon\right) + c^e(\Omega_{kt};\theta)}{\sigma_u}\right), 1\right\}$.

$$\left(\widehat{g^e(\Omega_{kt};\theta)} - \overline{g^e(\Omega_{kt})}\right) \cdot n^{not\_yet\_e}(\Omega_{kt})$$
$$\left(\widehat{g^d(\Omega_{kt};\theta)} - \overline{g^d(\Omega_{kt})}\right) \cdot n^{e\_not\_yet\_d}(\Omega_{kt})$$
$$\left(\left(E\left[\pi^d(\Omega_{kt},\varepsilon_{it};\theta)||I^d_{ikt}=1\right] - \overline{\pi^d(\Omega_{kt})}\right) \cdot n^d(\Omega_{kt})\right)$$
$$\Omega'_{kt}\left(\left(\widehat{g^e(\Omega_{kt};\theta)} - \overline{g^e(\Omega_{kt})}\right) \cdot n^{not\_yet\_e}(\Omega_{kt})\right)$$
$$\Omega'_{kt}\left(\left(\widehat{g^d(\Omega_{kt};\theta)} - \overline{g^d(\Omega_{kt})}\right) \cdot n^{e\_not\_yet\_d}(\Omega_{kt})\right)$$

where, for each state of the market Ω_{kt}, $n^{not\_yet\_e}(\Omega_{kt})$ is the number tracts that have yet to be explored, $n^{e\_not\_yet\_d}(\Omega_{kt})$ is the number of tracts that have been explored but not developed, and $n^d(\Omega_{kt})$ is the number of tracts that have been developed. The GMM estimator $\widehat{\theta}$ is then the solution to the problem:

$$\min_{\theta}\left[\frac{1}{n}\sum_{k,t}\Psi(\Omega_{kt},\theta)\right]' W_n^{-1}\left[\frac{1}{n}\sum_{k,t}\Psi(\Omega_{kt},\theta)\right], \tag{22}$$

where n is the number of observations (i.e., the number of market-time pairs) and W_n is a weight matrix.[23] In the present estimation, the number of moments that involved interactions with state variables is chosen so that the number of moments is equal to the number of parameters. Because the system is therefore exactly identified, the weight matrix used is the identity matrix.[24]

Identification of the parameters in the profit function comes from the realized profit function averages and in particular the moments that match, for those tracts that developed, the expected development profits conditional on development predicted by the model with the actual average realized profits in the data. Identification of the parameters $(\sigma_\mu, \sigma_\varepsilon)$ governing the distribution of private information is similar to the identification of the entry and exit parameters in Pakes et al. (2007): it comes

[23] The minimization algorithm used was a BFGS Quasi-Newton method with a cubic line search procedure. A grid search method was also tried, but this method took prohibitively long and tended not to result in weighted moments that were much lower than what resulted from the BFGS Quasi-Newton method with a cubic line search procedure.

[24] If the system is overidentified, a two-step GMM estimator can be used (Hansen, 1982; Graham, 2005). In the first step, a preliminary estimate of θ is obtained using the identity matrix as the weight matrix. In the second step, this preliminary estimate of θ is used to construct the optimal weighting matrix as specified by Chamberlain (1987), which is then used to obtain the final estimate of θ.

from the realized exploration and development frequencies and in particular the moments that match the predicted exploration and development probabilities with the actual probabilities in the data.

Standard errors are formed by a nonparametric bootstrap. Markets are randomly drawn from the data set with replacement to generate 100 independent panels of size equal to the actual sample size. The structural econometric model is run on each of the new panels. The standard error is then formed by taking the standard deviation of the estimates from each of the random samples.

The online Appendix presents results of Monte Carlo experiments that were run to assess the finite sample distribution of the estimators.

5 Results

Before estimating the structural model, I first run preliminary probits of the exploration decision and development decision, respectively. These models ignore the dynamic nature of the production decisions and do not address the endogeneity of neighbors' decisions, and thus may yield misleading results. The tract-year observations used for the exploration probit consist of tracts that have not yet been explored before time t. The tract-year observations used for the development probit consist of tracts that have been explored but have not yet been developed before time t. The results for the exploration and development probits by tract size are reported in Tables 4 and 5, respectively. According to the results, when dynamic and endogeneity considerations are ignored, neither the marginal effect of the number of tracts in the market that have been explored nor the marginal effect of the number of tracts in the market that have been developed is statistically significant either for the entire pooled sample or for the subsample of large tracts, or for the subsample of small tracts. Thus, according to these results, there are no strategic interactions in either stage of petroleum production, even on the small tracts. Because these probits ignore dynamic and endogeneity considerations, however, their results are less realistic than the results of the structural model that follow and thus may be misleading indicators of the actual strategic interactions that take place.

A feature of the data that suggests that there may be externalities and therefore that the reduced-form probits may be giving a misleading result can be seen in Figure 5, which plots the aggregate hazard rate of exploration. The aggregate hazard rate H_t

at time t is computed as the number of tracts that explored at time t divided by the risk set R_t at time t, where the risk set is simply the set of tracts that have not explored before time t. Following Hendricks and Porter (1996), the standard deviation of the hazard is $\sqrt{H_t (1 - H_t) / R_t}$; the error bars in the figure indicate plus or minus one standard deviation. As seen in this figure, the aggregate hazard rate exhibits a U-shaped pattern: the hazard rate of exploration is monotonically decreasing with time except in the year right before the lease term expires, when there is a spike in exploration. Since firms are subject to a lease term by the end of which they must begin exploratory drilling, or else relinquish their lease, then this U-shaped hazard rate is consistent with the presence of an information externality, which would result in too little exploration at the beginning of the lease term and duplicative drilling in the final period of the lease (Hendricks & Porter, 1996; Porter, 1995). This result is therefore suggestive that an information externality may be present that leads firm to play a non-cooperative timing game that leads them to inefficiently delay production, since the possibility of acquiring information from other firms may further enhance the option value to waiting.

For the structural model, there are three different types of parameters to be estimated: the parameters $(\sigma_\mu, \sigma_\varepsilon)$ governing the distribution of the private information, the coefficient α on drilling costs in the exploration profit function, and the coefficients γ on the state variables in the development profit function.

The parameters estimated in the structural model include the coefficients in the payoff function from developing a tract. The payoff function itself is specified in a reduced-form fashion since it is assumed to be linear in the state variables. However, the payoff function is imbedded in a multi-stage, dynamic investment timing game model, and the parameters in the payoff function are estimated by matching the investment policy functions obtained from solving the multi-stage dynamic investment timing game model to the actual investment probabilities in the data. Thus, while the way the parameters enter into the payoff function is reduced-form, the way the payoff function (and thus the parameters in the payoff function) enters into the investment policy function is structural, and based on the structural multi-stage dynamic investment timing game model. Unlike the static reduced-form discrete response models used by the previous literature and the preliminary probits estimated above, where the investment probability is specified as a reduced-form function of the regressors, in this structural estimation the model for the investment probability is based on a solution to a multi-stage dynamic game, and is therefore a function not only of the parameters in the payoff function, but also of the continuation value to waiting and the parameter in the distribution of private information. Thus, the investment policy function derived in this paper accounts for the option value to waiting before making an irreversible investment, and thus is different from a static investment rule that would not take into account the option value to waiting. The option value to waiting is thus essentially an omitted variable in the static reduced-form models used in the previous literature, and since this option value to waiting is likely correlated with the regressors, there is omitted variable bias in the static reduced-form models. While the parameters estimated in this current model include those that can be directed interpreted as measuring the relationship between various state variables and the profits of firms, albeit a reduced-form relationship, the parameters in the discrete response regressions of investment in the previous literature are not directly interpretable as relating to profits, suffer from omitted variables bias, and confound static and dynamic considerations.

The coefficients on the strategic variables in this paper can be interpreted as the net effect of a neighbor's action on a firm's payoff. In contrast, it is unclear how to interpret the strategic parameter estimated in the reduced-form response models of previous papers in terms of its effects on profits. Since the parameters in this paper directly relate to the payoff function rather than to a reduced-form specification of investment probability, in this paper I am able to run welfare analyses to evaluate the inefficiencies arising from strategic interactions relative to what a firm that owned both tracts in the neighborhood and therefore internalized the externalities would do.

The results from running the structural model on all tracts in the panel regardless of tract size are shown in Table 6. Two specifications are reported: the base case specification (1) and an alternative specification (2) where dummies for the medium and high oil price bins are used instead of one continuous variable for oil price bin. The results for specification (2) are similar to the results for specification (1), and the coefficients on the dummy variables for the oil price bins in specification (2) are similar to each other. The discussion thus focuses on the base case specification (1). The coefficients $\gamma_N \equiv (\gamma_{tote}, \gamma_{totd})$ on the two endogenous state variables can be interpreted as follows. Since the total number of tracts in the market that have been explored increases by one if a neighbor explores, the coefficient γ_{tote} on this variable measures how the profits from development change when a neighbor explores. Similarly, since the total number of tracts in the market that have been developed can only take values of 0 and 1, with 0 indicating that the neighbor has not developed and 1 indicating that the neighbor has developed, the coefficient γ_{totd} on this variable measures how the profits from development change when a neighbor develops first. An important result is that the coefficient on the total number of tracts in the market that have been explored is statistically insignificant, which means that firms do not interact strategically on net during exploration. In contrast, the coefficient on the total number of tracts in the market that have been developed is statistically significant and positive, which means that a firm's profits increase when its neighbor develops. This seems reasonable, because when a neighbor develops following exploration, this is a signal to the firm that the neighbor's exploratory efforts were successful, and therefore that there may be deposits present.

To assess the economic significance of the strategic interactions, I compare the coefficients γ_N on the endogenous state variables with the development cost $|\gamma_{drill} \cdot drill\_cost_t|$, where γ_{drill} denotes the coefficient on the discretized real drilling cost $drill\_cost_t$ in the development profit function. As noted above, one would expect each coefficient γ_N to be less than the cost of development. Since the maximum development cost is given by $|\gamma_{drill} \cdot \max(drill\_cost_t)| = |\gamma_{drill}|$, one would therefore expect $\gamma_{tote} < |\gamma_{drill}|$ and $\gamma_{totd} < |\gamma_{drill}|$. The results are consistent with the expectations. Further comparison of the coefficients γ_N with the mean development cost $\left|\gamma_{drill} \cdot \overline{drill\_cost_t}\right|$ can give a measure of the economic importance of a neighbor's decisions to a firm's profits. In particular, the relative importance of a neighbor's exploration decision as a fraction of a firm's costs is given by $\frac{\gamma_{tote}}{\left|\gamma_{drill} \cdot \overline{drill\_cost_t}\right|} = \frac{\gamma_{tote}}{0.49|\gamma_{drill}|} = 0.01$ in specification (1). Similarly, the relative importance of a neighbor's development decision as a fraction of a firm's costs is given by $\frac{\gamma_{totd}}{\left|\gamma_{drill} \cdot \overline{drill\_cost_t}\right|} = \frac{\gamma_{totd}}{0.49|\gamma_{drill}|} = 0.03$ in specification (1). The small values of both of these numbers indicate that the effects of neighbors' decisions are second-order compared to costs.

In absolute terms, strategic interactions in exploration are not only statistically insignificant, but economically insignificant as well: according to specification (1), decreases in profits (in 1982 $) more than $10,000 and increases in profits more than $70,000 resulting from a neighbor's exploration can be rejected at a 5% level. These values are small relative to predicted ex post revenues, which average $49.34 million. Strategic interactions in development are statistically significant but only moderately economically significant: increases in profits resulting from a neighbor's development that are less than \$120,000 and greater than \$200,000 can be rejected at a 5% level.

The lack of strategic effects during exploration is consistent with Lin (2009), who estimates a reduced-form discrete response model of a firm's exploration timing decision that uses variables based on the timing of a neighbor's lease term as instruments for the neighbor's decision, and whose results show that information and extraction externalities do not appear to induce firms to interact strategically on net during exploration.

As for the values of the parameters governing the distribution of private information, both the parameter σ_μ from the distribution of the pre-exploration shock μ_{ikt} and the parameter σ_ε from the distribution of the pre-development shock ε_{ikt} are statistically significant. This suggests that both the mean and the variance of the private information have a statistically significant impact on both the exploration decision and the development decision.

In terms of economic significance, one way to interpret the mean σ_μ of the pre-exploration shock μ_{it} is to compare it with exploration costs. Since both the pre-exploration shock μ_{ikt} and the exploration cost function $c^e(\Omega_{kt};\theta)$ enter linearly into the exploration profit function (11), the importance of private information in the exploration decision can be measured by comparing the mean σ_μ of the shock with the mean $\overline{c^e(\Omega_{kt};\theta)}$ of the costs. Expressed as a fraction of the mean exploration costs, where the mean exploration costs are computed by substituting the mean value of the discretized real drilling cost into the equation (9) for the exploration costs, the relative importance of private information is therefore given by: $\frac{\sigma_\mu}{\overline{c^e(\Omega_{kt};\theta)}} = \frac{\sigma_\mu}{-\alpha\cdot\left(\overline{drill\_cost_t}+1\right)} = \frac{\sigma_\mu}{-1.49\alpha}$. A large value of $\frac{\sigma_\mu}{-1.49\alpha}$ would indicate that private information plays a large role relative to costs in the first-stage exploration decision; a small value would indicate that costs are more relatively more important. In this case, the value is 0.33. Private information is about a third as important as costs. Thus, the role of private information in the exploration decision is both economically and statistically significant.

The mean σ_ε of the pre-development shock ε_{ikt} can be similarly compared with development costs to assess the importance of private information in the second-stage development decision. Since both the pre-development shock ε_{ikt} and the development cost $|\gamma_{drill} \cdot drill\_cost_t|$ enter linearly into the development profit function (1), the importance of private information in the development decision can be measured by comparing the mean σ_ε of the shock with the mean $\left|\gamma_{drill} \cdot \overline{drill\_cost_t}\right|$ of the costs. Expressed as a fraction of the mean development costs, the relative importance of private information is therefore given by: $\frac{\sigma_\varepsilon}{\left|\gamma_{drill} \cdot \overline{drill\_cost_t}\right|} = \frac{\sigma_\varepsilon}{0.49|\gamma_{drill}|}$. A large value of $\frac{\sigma_\varepsilon}{0.49|\gamma_{drill}|}$ would indicate that private information plays a large role relative to costs in the second-stage development decision; a small value would indicate that costs are more relatively more important. In this case, the value is 1.00 to 1.02. Private information is approximately as important as costs in the development decision.

The parameter σ_μ from the distribution of the pre-exploration shock μ_{ikt} and the parameter σ_ε from the distribution of the pre-development shock ε_{ikt} are measures not only of the mean of the respective distributions, but also of their variances as well, which are, respectively, σ_μ^2 and σ_ε^2. The statistical significance of both these parameters suggests that, in addition to the mean, the variance in the private information has a statistically significant impact on both the exploration decision and the development decision. The relative importance $\frac{\sigma_\mu^2}{-1.49\alpha}$ of the variance in the private information relative to the costs in the first-stage exploration decision is 1.67. The relative importance $\frac{\sigma_\varepsilon^2}{0.49|\gamma_{drill}|}$ of the variance in the private information relative to the costs in the second-stage development decision is 2.74 to 5.04. Thus, the variance in the private information is both economically and statistically significant.

The coefficients on the other covariates are all significant and have the expected sign. As expected, the coefficient on the discretized drilling cost is negative in both the exploration profit equation and the development profit equation. Also as expected, the profits from development increase in both the average winning bid and in the real oil price.

There are several possible explanations why the results do not provide evidence for strategic, non-cooperative behavior during exploration. One is that the tract sizes are large enough that cross-tract externalities are insignificant or even nonexistent. A second is that cross-tract externalities exist, but firms owning neighboring tracts cooperate to jointly internalize the externalities they impose on each other, for example through joint ventures or unitization. A third is that cross-tract externalities are significant, but the positive information externality exactly cancels the negative extraction externality, resulting in zero net strategic effects.

To distinguish among these three explanations for the lack of strategic interactions during exploration, I estimate the strategic interactions by tract size. Due to the smaller sample size, and because the results for the two specifications are similar, the estimation of the model by tract size is done on the base case specification (1) only.

If externalities are insignificant when tracts are large enough, then one would expect to see strategic, non-cooperative behavior only on small tracts. This is because the smaller the tract size, the more likely the tracts are located over a common pool, and therefore the more acute the information and extraction externalities faced by the firms. Evidence for significant strategic interactions on small tracts but not on large tracts would thus be consistent with the first explanation.

If externalities exist even for the largest tracts in the sample, but are eliminated through coordination, then, assuming firms would coordinate regardless of tract size, one would not expect to see any strategic, non-cooperative behavior even on small tracts. Insignificant strategic interactions regardless of tract size would thus be consistent with the second explanation.[25] Neighboring tracts might coordinate by forming joint ventures in exploration or by consolidating their production rights through purchase or unitization. Joint ventures in exploration occur less frequently than one might expect, however, because negotiations are contentious, because firms fear allegations of pre-sale anti-trust violations, and because prospective partners have an incentive to free ride on a firm's information gathering expenditures (Hendricks and Porter, 1992). In their theoretical model of the persuasion game, Hendricks and Kovenock (1989) find that, even with well-defined property rights, bargaining does not eliminate all the inefficiencies of decentralized drilling decisions. As a consequence, the information externality may not be fully internalized. Under a unitization agreement, a single firm is designated as the unit operator to develop the entire reservoir, while the other firms share in the profits according to negotiated formulas (Libecap & Smith, 1999). There are many obstacles to consolidation, however, including contentious negotiations, the need to determine relative or absolute tract values, information costs, and oil migration problems (Libecap & Wiggins, 1984). In addition, another free rider problem that impedes coordination is that firms may fear that if they reveal to other firms their information or expertise, for example about how to interpret seismic data, then they may lose their advantage in future auctions (Hendricks & Porter, 1996). Thus, despite various means of coordination, firms may still behave strategically and non-cooperatively, and information and extraction

[25] It is also possible that the coefficients that arise when firms coordinate are significant, but are different from those that would arise under the non-cooperative outcome.

externalities may not be fully internalized.[26]

If strategic interactions do not occur on net because the positive information externality exactly cancels the negative extraction externality, then one may not expect the exact cancellation to still take place when the tract size is small. This is because the geographical span of the information externality is larger than that of the extraction externality: while the former only requires that tracts may share common geological features, the latter requires that tracts may be located over a common pool. If neighboring tracts share common geological features, then they can benefit from the information about these geological features gleaned from each other's investment timing decisions, but neighboring tracts will only compete for the same pool of oil if they are located over a common pool. Because the geographical span of an area sharing common geological features is generally larger than the geographical span of a pool of petroleum (Divi, 2004), the geographical span of the information externality is larger than that of the extraction externality. As a consequence, it is possible for the information externality to be present on all the tracts in the sample, but for the extraction externality to be present on only the smaller tracts. Theory therefore suggests that the importance of the extraction externality relative to the information externality should be greater on small tracts than on large tracts. Thus, if the externalities cancel when all the tract sizes are considered, one might expect that the negative extraction externality would dominate the positive information externality when the sample is limited to small tracts only. Strategic interactions that are more significantly negative on small tracts than on large tracts would therefore be consistent with the third explanation.

Running the structural model on subsamples of the data set that differed in the acreage of the tracts in the panel would therefore enable one to distinguish among these three explanations for the lack of strategic interactions in the pooled sample. Assuming that the tract sizes differ for exogenous reasons, the results by acreage will also give a sense of whether or not the government can change the extent the which firms behave strategically by changing the size of the tracts.

---

[26] In addition to cooperating during exploration, firms can also cooperate in development via unitization. However, "unitization is not nearly as prevalent in the OCS as it is in onshore areas. Some of the more obvious reasons for this are there is only one lessor, the federal government, leases are larger, usually 5,000 or 5,760 acres and in most cases, the royalty is the same on adjacent blocks" (Newton, 1994, p.21). Moreover, the significant positive strategic effect during development in the pooled sample is further corroboration that cooperation during development does that occur as much as one might expect.

Table 7 presents the results from running the structural model on a subsample consisting of larger tracts, defined as tracts that are greater than or equal to 5000 acres in size. As the coefficients on the endogenous variables indicate, strategic interactions are neither economically nor statistically significant. For tracts greater than or equal to 5000 acres in size, the 95% confidence interval for the effect of a neighbor's exploration on profits (in 1982 $) ranges from -$8.48 million to $8.48 million, and the 95% confidence interval for the effect of a neighbor's development on profits ranges from -$2.24 million to $2.24 million. These values are small relative to predicted ex post revenues, which average $49.34 million in the pooled sample. Investment decisions and profits are instead driven primarily by private information, the average winning bid, the real drilling cost, and the real oil price.

Table 8 presents the results from running the structural model on three subsamples consisting of smaller tracts, defined as tracts that are less than 5000 acres, 4000 acres, and 3000 acres in size, respectively. Strategic interactions are statistically and economically significant in all three subsamples. The coefficients in the three specifications on the number of tracts in the market that have been explored indicate that real development profits decrease by a statistically significant $25.78 million to a statistically significant $27.67 million when a neighbor explores; a neighbor's exploration is roughly as important to profits as maximum development costs. The negative extraction externality thus appears to dominate during exploration: when a neighbor explores, a firm's profits decrease because the neighbor has begun production and is likely to eventually compete with the firm for the same common pool. The coefficients in the three specifications on the number of tracts in the market that have been developed indicate that real development profits increase by a statistically significant $3.15 million to a statistically significant $4.06 million when a neighbor develops. The positive information externality dominates during development: a firm benefits when its neighbor develops after it explores because this is a signal to the firm that the neighbor's exploratory efforts were successful, and therefore that there may be deposits present. The magnitude of the positive net strategic effect that results from a neighbor's development is one order of magnitude smaller than the negative net strategic effect that results from a neighbor's exploration. As expected, the magnitude of the strategic interactions in both exploration and development increase monotonically as the tract sizes get smaller.

For the smaller tract sizes, investment decisions and profits are also driven by private information, the average winning bid, the real drilling cost, and the real oil price, as before. Pre-development private information plays a larger role in decision-making and profits for the smaller tracts than it does for the larger tracts.

The results therefore indicate that there are no strategic effects on larger tracts because the tract size is large enough to enable a firm to internalize the externalities on its own. On small tracts, however, externalities are acute and the positive information externality dominates. Small tracts make up a small percentage of the tracts sold: of the tracts used in the structural estimation, only 37%, 27%, and 24% of tracts were less than 5000 acres, 4000 acres and 3000 acres in size, respectively. Thus, for the majority of tracts, externalities do not cause inefficient strategic interactions on net. The results of the estimations by tract size show that the reason there is a lack of strategic interactions in the pooled sample is that the tract sizes are large enough that cross-tract externalities are insignificant or even nonexistent, and not that neighboring tracts cooperate or that the information and extraction externalities cancel each other out. The results suggest that, by making most of the tracts at least 5000 acres in size, the federal government has minimized the net effects of any externalities that may be present, and has thus avoided additional inefficiencies in petroleum production that may have resulted from non-cooperative strategic interactions.

In their analysis of drilling decisions, Hendricks and Porter (1996) find that some evidence that equilibrium predictions of plausible non-cooperative models are reasonably accurate and more descriptive than those of cooperative models of drilling timing. The results of their probit of the probability of initial drilling by year after acquisition show that the coefficient on (the logarithm of) the total number of tracts explored since the sale date is borderline significant and positive two years after acquisition, but not significant for any of the other years during the lease term. They also find that the coefficient on (the logarithm of) the number of drilled tracts that were productive is significant and positive three years after acquisition, but not significant for any of the other years during the lease term.

Although Hendricks and Porter (1996) do find some strategic interactions during exploration, their evidence for strategic interactions is somewhat weak, as the relevant coefficients are not always statistically significant. This paper refines the Hendricks and Porter (1996) analysis by analyzing the development decision in addition to the exploratory drilling decision; by estimating a structural model of a dynamic game that incorporates the dynamic, multi-stage nature of petroleum production and that allows for a possible option value to waiting before either exploring or developing; by allowing for extraction externalities as well as information externalities that arise during both exploration and development; and by estimating the strategic interactions by tract size. I find no evidence of strategic interactions on large tracts and significant strategic interactions on small tracts. I also find a significant positive net effect from a neighbor's development in both the pooled sample and on the small tracts, which is evidence that the information externality is important in the second stage of petroleum production, for when a neighbor develops following exploration, this is a signal to the firm that the neighbor's exploratory efforts were successful, and therefore that there may be deposits present. The lack of strategic interactions during exploration for the pooled sample, the strong evidence for strategic interactions in both stages of production on the small tracts, and the positive effect of a neighbor's development both in the pooled sample and for the small tracts are all results of this paper that differ from and are refinements to Hendricks and Porter's (1996) results, results that differ due to the improvements made by this paper on their seminal work.

To analyze the effects of the information and extraction externalities on tract exploration, tract development, and profits from development, the estimated distributions of the parameters were used to simulate the policies that would be chosen if a firm owned both tracts in the market rather than just one (or if the two firms colluded or cooperated perfectly with each other). For each market, and for each possible combination of exploration and development policies for both tracts, the joint expected profits were determined by averaging the profits from simulating 100 paths for the exogenous variables using the initial conditions for the exogenous variables and the empirical transition matrix for the exogenous variables. The exploration and development policies for both tracts that maximized the joint expected profits for both tracts in the market were then determined for each market, then averaged over all markets. To account for the standard errors in the parameter estimates, results are averaged over all markets for each of 100 draws from the estimated distribution of parameters. The simulations were run for all tracts used in the pooled analysis as well as for tracts less than 5000 acres, 4000 acres and 3000 acres, respectively. The results are shown in Table 9. According to the results, the externalities cost between 16.51 and 25.31 million real US dollars (1982 \$) per tract developed. In the data, when two neighboring tracts that are not adjacent to any other tracts are owned by the same firm rather than different firms, results from a regression of the profits on a dummy for the tracts being owned by the same firm show that tracts owned by the same firm earn a statistically significant 2.08 million real US dollars (1982 $) more than tracts owned by different firms. One reason the externality cost in the data is lower than the externality cost calculated by the simulations is that, in the data, tracts that are owned by the same firm may be a select sample, and may differ from tracts that are owned by different firms. However, the data still show that joint ownership yields more profits than separate ownership does.

6 Conclusion

When individual petroleum-producing firms make their exploration and development investment timing decisions, information externalities and extraction externalities may lead them to interact strategically with their neighbors. A positive information externality arises if tracts are located over a common pool or share common geological features so that their ex post values are correlated, since firms learn information about their own tracts when other firms drill exploratory wells or install production platforms on neighboring tracts. A negative extraction externality arises when tracts are located over a common pool, since firms are competing for the same stock of petroleum. Owing to both information and extraction externalities, the dynamic decision-making problem faced by a petroleum-producing firm is not merely a single-agent problem, but rather can be viewed as a multi-agent, non-cooperative game in which firms behave strategically and base their exploration and development policies on those of their neighbors.

This paper examines whether strategic considerations arising from information and extraction externalities are present. In particular, it analyzes whether a firm's investment timing decisions and profits depend on the decisions of firms owning neighboring tracts of land. The econometric approach employed is a structural econometric model of the firms' multi-stage investment timing game.

The research presented in this paper is important for several reasons. First, an empirical analysis of investment timing decisions enables one to examine whether the strategic interactions that are predicted in theory actually occur in practice. Second, the estimation of strategic interactions, especially those that arise in dynamic decision-making, is of methodological interest. The structural econometric methodology employed can be used to analyze externalities in a variety of contexts, including spillovers that arise during research and development. Third, my results have implications for leasing policy: if the strategic effects and externalities turn out to be large, then the program by which the U.S. government leases tracts to firms may be inefficient, and possible modifications should be considered.

Do the positive information externalities and negative extraction externalities have any net strategic effect? The answer depends on tract size. When the tract sizes are large, firms do not impose externalities on each other on net when choosing to explore or develop, and, as a consequence, strategic considerations are second-order. This is the case with most of the tracts in the federal leasing program. However, in the few cases where the tract size is small, externalities do matter, and they cause firms to interact strategically with their neighbors. As expected, these externalities intensify as the tract size decreases. Also as predicted by theory, during exploration the relative importance of the extraction externality with respect to the information externality is greater on small tracts than on large tracts. The results suggest that, in making most tracts at least 5000 acres in size, the federal government has avoided additional inefficiencies in petroleum production that may have resulted from non-cooperative strategic interactions.

References

[1] Aguirregabiria, V., & Mira, P. (2007). Sequential estimation of dynamic discrete games. Econometrica, 75 (1), 1-53.

[2] Bajari, P., Benkard, C.L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. Econometrica, 75, 1331-1370.

[3] Casella, G., & Berger, R. (1990). Statistical inference. Belmont, CA: Duxbury Press.

[4] Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. Journal of Econometrics, 34, 305-334.

[5] Chermak, J.M., Crafton, J., Norquist, S.M., & Patrick, R.H. (1999). A hybrid economic-engineering model for natural gas production. Energy Economics, 21, 67-94.

[6] Divi, R.S. (2004). Probabilistic methods in petroleum resource assessment, with some examples using data from the Arabian region. Journal of Petroleum Science and Engineering, 42, 95-106.

[7] Dixit, A., & Pindyck, R. (1994). Investment under uncertainty. Princeton, NJ: Princeton University Press.

[8] Fudenberg, D., & Tirole, J. (1998). Game theory. Cambridge: MIT Press.

[9] Graham, B. (2005). Small sample properties of GMM and GEL estimation and inference procedures for social interaction models. Working paper. University of California at Berkeley.

[10] Hansen, L. (1982). Large sample properties of generalized method of moments estimators. Econometrica, 50 (4), 1029-1054.

[11] Harsanyi, J. (1973). Games with randomly disturbed payoffs. International Journal of Game Theory, 2, 1-23.

[12] Hendricks, K., & Kovenock, D. (1989). Asymmetric information, information externalities, and efficiency: The case of oil exploration. RAND Journal of Economics, 20 (2), 164-182.

[13] Hendricks, K., & Porter, R. (1992). Joint bidding in federal OCS auctions. American Economic Review Papers and Proceedings, 82 (2), 506-511.

[14] Hendricks, K., & Porter, R. (1993). Determinants of the timing and incidence of exploratory drilling on offshore wildcat tracts. NBER Working Paper Series (Working Paper No. 4605). Cambridge, MA.

[15] Hendricks, K., & Porter, R. (1996). The timing and incidence of exploratory drilling on offshore wildcat tracts. The American Economic Review, 86 (3), 388-407.

[16] Hendricks, K., Porter, R., & Boudreau, B. (1987). Information, returns, and bidding behavior in OCS auctions: 1954-1969. Journal of Industrial Economics, 35 (4), 517-542.

[17] Hurn, A., & Wright, R. (1994). Geology or economics? Testing models of irreversible investment using North Sea oil data. Economic Journal, 104 (423), 363-371.

[18] Kellogg, R. (2010). The effect of uncertainty on investment: Evidence from Texas oil drilling. NBER Working Paper Series (Working Paper No. 16541). Cambridge, MA.

[19] Libecap, G., & Smith, J. (1999). The self-enforcing provisions of oil and gas unit operating agreements: Theory and evidence. The Journal of Law, Economics, and Organizations, 15 (2), 526-548.

[20] Libecap, G., & Smith, J. (2001). Regulatory remedies to the common pool: The limits to oil field unitization. Energy Journal, 22 (1), 1-26.

[21] Libecap, G., & Wiggins, S. (1984). Contractual responses to the common pool: Prorationing of crude oil production. The American Economic Review, 74 (1), 87-98.

[22] Libecap, G., & Wiggins, S. (1985). The influence of private contractual failure on regulation: The case of oil field unitization. The Journal of Political Economy, 93 (4), 690-714.

[23] Lin, C.-Y.C. (2009). Estimating strategic interactions in petroleum exploration. Energy Economics, 31 (4), 586-594.

[24] Manski, C. (1993). Identification of endogenous social effects: The reflection problem. Review of Economic Studies, 60 (3), 531-542.

[25] Manski, C. (1995). Identification problems in the social sciences. Cambridge, MA: Harvard University Press.

[26] Newton, J.A. (1994). The Outer Continental Shelf: A Home Study Course for Continuing Professional Development in Land. Fort Worth, TX: American Association of Professional Landmen.

[27] Paddock, J., Siegel, D., & Smith, J. (1988). Option valuation of claims on real assets: The case of offshore petroleum leases. Quarterly Journal of Economics, 103 (3), 479-508.

[28] Pakes, A., Ostrovsky, M., & Berry, S. (2007). Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples). RAND Journal of Economics, 38(2), 373-399.

[29] Pesaran, M.H. (1990). An econometric analysis of exploration and extraction of oil in the U.K. Continental Shelf. The Economic Journal, 100 (401), 367-390.

[30] Porter, R. (1995). The role of information in U.S. offshore oil and gas lease auctions. Econometrica, 63 (1), 1-27.

[31] Robalino, J., & Pfaff, A. (2005). Contagious development: Neighbors' interactions in deforestation. Working paper. Columbia University.

[32] Starr, A.W., & Ho, Y.C. (1969). Nonzero-sum differential games. Journal of Optimization Theory and Applications, 3, 184-206.

[33] Stokey, N., Lucas, R., & Prescott, E. (1989). Recursive methods in economic dynamics. Cambridge, MA: Harvard University Press.

FIGURE 1.

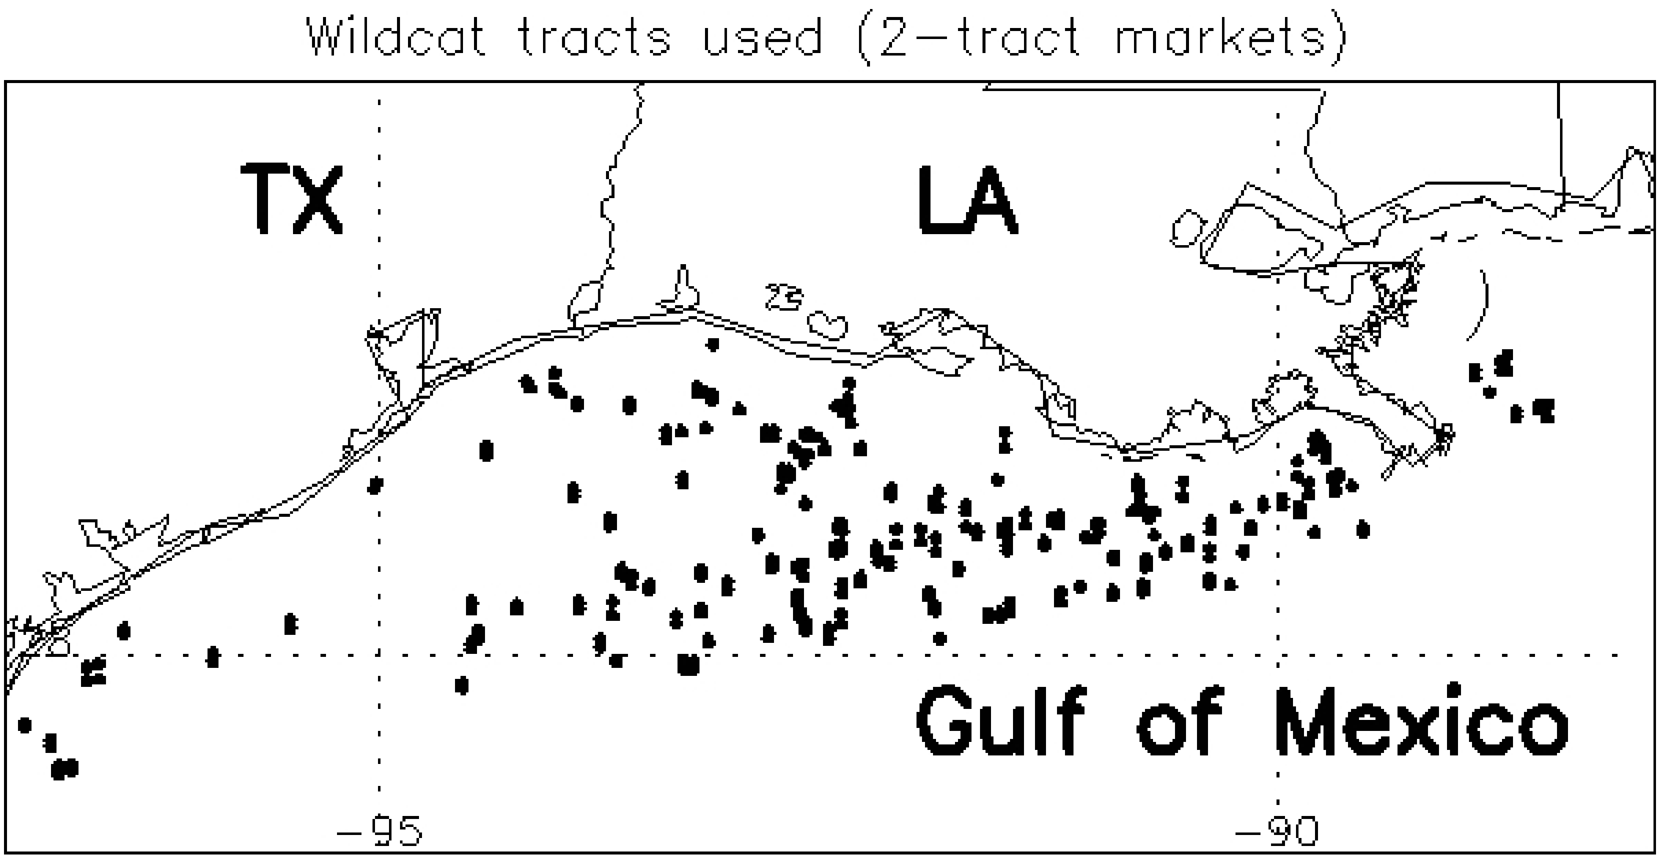
Wildcat tracts used (2-tract markets)
TX
LA
Gulf of Mexico
-95
-90

TABLE 1. Summary statistics for wildcat tracts used in analysis (2-tract markets) and tracts not used in analysis

	Tracts used	Tracts not used
number of years to exploration, conditional on exploring	1.21	1.27
	(1.42)	(1.41)
number of years to development, conditional on developing	5.79	5.38
	(3.51)	(3.18)
revenue (million 1982 $), conditional on developing	49.34	48.80
	(65.53)	(82.6)
gross profits (million 1982 $), conditional on developing	-10.83	-8.76
	(9.57)	(11.1)
# tracts	174	4860

FIGURE 2.

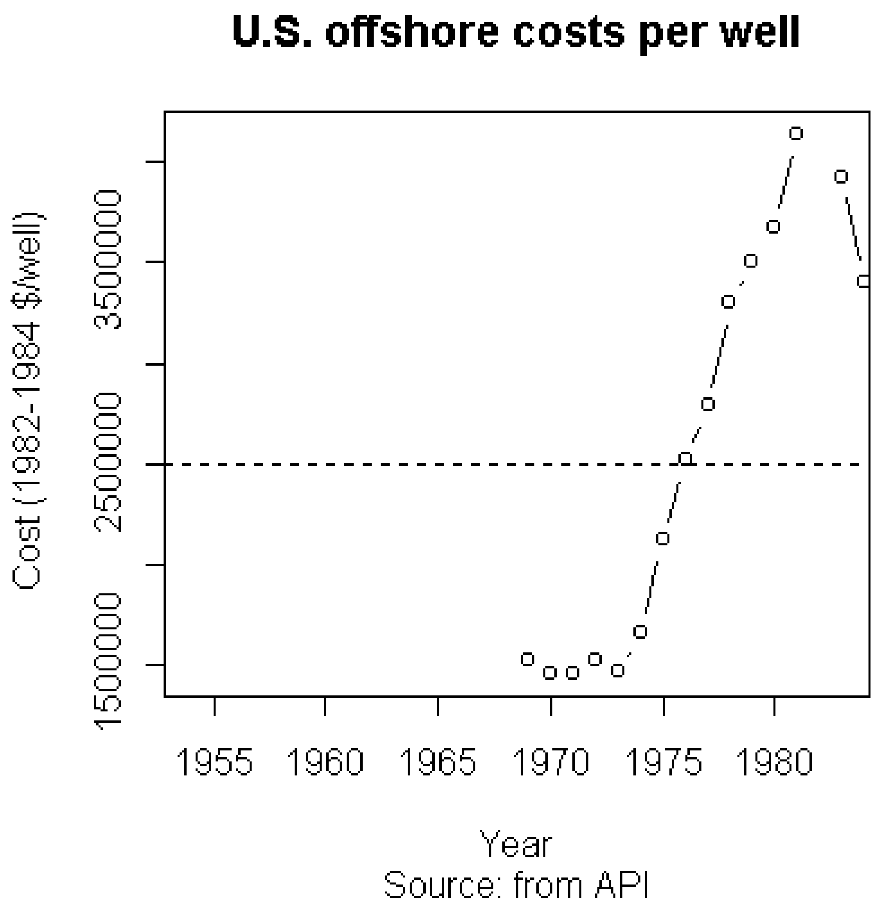
U.S. offshore costs per well
Cost (1982-1984 $/well)
1500000
2500000
3500000
1955
1960
1965
1970
1975
1980
Year
Source: from API

FIGURE 3.

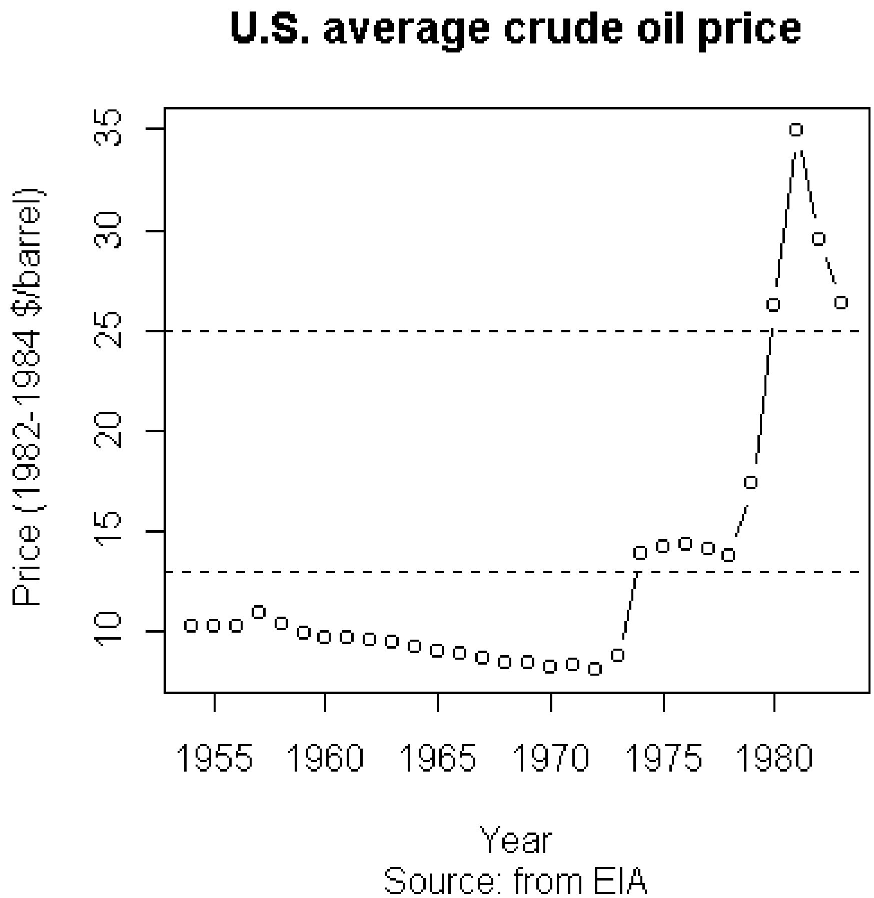
U.S. average crude oil price
Price (1982-1984 $/barrel)
10
15
20
25
30
35
1955
1960
1965
1970
1975
1980
Year
Source: from EIA

TABLE 2. Summary statistics

	# obs	mean	s.d.	min	max
by tract					
acreage (1000 acres)	174	4.46	1.30	0.94	5.76
number of years to exploration, conditional on exploring	122	1.21	1.42	0	4
number of years to development, conditional on developing	66	5.79	3.51	0	15
revenue (million 1982 $), conditional on developing	66	49.34	65.53	0.022	298.0
gross profits (million 1982 $), conditional on developing	66	-10.83	9.57	-38.10	18.80
by market-year					
# tracts in market that have been explored	1041	1.49	0.80	0	2
# tracts in market that have been developed	1041	0.31	0.46	0	1
discretized average winning bid per acre	1041	0.76	0.65	0	2
discretized real drilling cost	1041	0.42	0.49	0	1
discretized real oil price	1041	0.57	0.69	0	2

FIGURE 4.

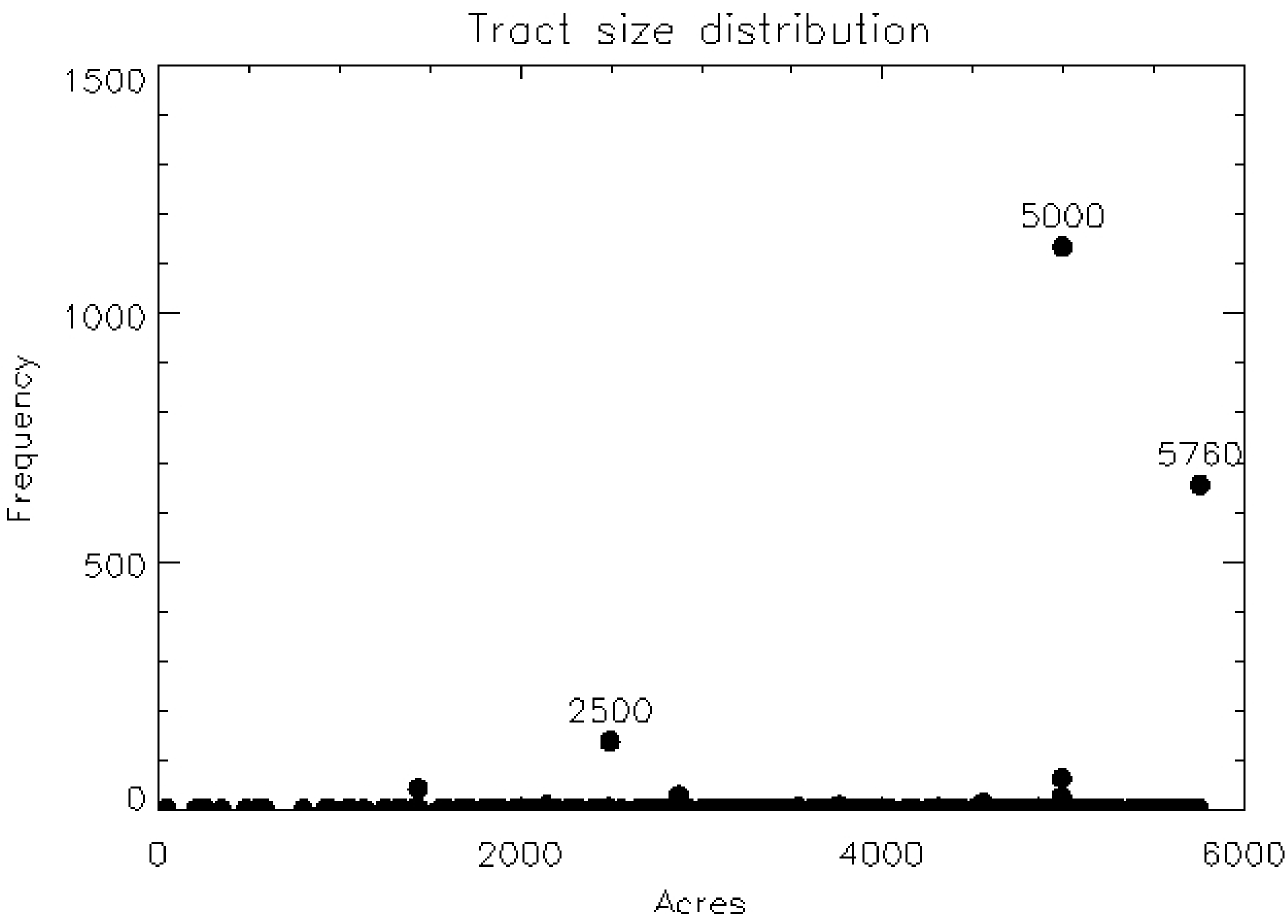
Tract size distribution
1500
1000
500
0
Frequency
5000
5760
2500
0
2000
4000
6000
Acres

TABLE 3. Summary statistics by tract size

	all tracts	Acreage ≥ 5000	< 5000	< 4000	< 3000
by tract					
fraction of tracts that are explored	0.70	0.75	0.60	0.60	0.55
fraction of tracts that are developed	0.38	0.42	0.29	0.29	0.24
number of years to exploration, conditional on exploring	1.21	0.98	1.74	2.24	2.38
	(1.42)	(1.24)	(1.75)	(1.69)	(1.75)
number of years to development, conditional on developing	5.79	5.72	5.35	5.83	6.78
	(3.51)	(3.10)	(4.51)	(4.90)	(5.29)
revenue (million 1982 $), conditional on developing	49.34	46.53	56.53	41.15	48.55
	(65.53)	(62.34)	(79.46)	(56.89)	(64.42)
gross profits (million 1982 $), conditional on developing	-10.83	-11.61	-7.27	-8.65	-9.08
	(9.57)	(8.77)	(9.62)	(7.85)	(9.12)
by market-year					
# tracts in market that have been explored	1.49	1.67	1.31	1.33	1.26
	(0.80)	(0.60)	(0.89)	(0.90)	(0.92)
# tracts in market that have been developed	0.31	0.20	0.27	0.30	0.21
	(0.46)	(0.40)	(0.44)	(0.46)	(0.41)
discretized average winning bid per acre	0.76	1.20	0.71	0.67	0.76
	(0.65)	(0.65)	(0.64)	(0.65)	(0.64)
discretized real drilling cost	0.42	0.48	0.33	0.33	0.30
	(0.49)	(0.51)	(0.47)	(0.47)	(0.46)
discretized real oil price	0.57	0.72	0.44	0.44	0.40
	(0.69)	(0.69)	(0.66)	(0.67)	(0.65)
# markets	87	55	29	21	19
# observations	1041	698	308	247	210

TABLE 4. Preliminary probits: Exploration

Dependent variable is dumming for exploring					
			Acreage		
	all tracts	≥ 5000	< 5000	< 4000	< 3000
other tract in market has been explored (dummy)	-0.080	-0.94	-0.07	0.05	0.02
	(0.04)	(0.06)	(0.06)	(0.10)	(0.09)
other tract in market has been developed (dummy)	0.26	0.34	0.25	0.13	0.05
	(0.16)	(0.22)	(0.27)	(0.23)	(0.20)
discretized average winning bid per acre	0.15 ***	0.14 ***	0.21 ***	0.16 ***	0.17 ***
	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
discretized real drilling cost	-0.003	-0.059	0.061	-0.048	-0.011
	(0.088)	(0.102)	(0.214)	(0.127)	(0.188)
discretized real oil price	-0.003	-0.041	0.012	0.12	0.060
	(0.058)	(0.074)	(0.123)	(0.18)	(0.189)
mean exploration probability	0.22	0.28	0.17	0.15	0.14
# observations	527	297	211	166	156

Notes: Marginal effects are reported. Standard errors in parentheses. Significance codes: * 5% level, ** 1% level, and *** 0.1% level.

TABLE 5. Preliminary probits: Development

Dependent variable is dummy for developing					
			Acreage		
	all tracts	≥ 5000	< 5000	< 4000	< 3000
other tract in market has been explored (dummy)	0.003	0.012	0.003	-0.009	0.010
	(0.009)	(0.011)	(0.018)	(0.018)	(0.020)
other tract in market has been developed (dummy)	0.004	-0.010	0.030	-0.003	-0.007
	(0.013)	(0.013)	(0.027)	(0.026)	(0.021)
discretized average winning bid per acre	0.029 ***	0.033 ***	0.030	0.019	0.020
	(0.007)	(0.009)	(0.014)	(0.015)	(0.014)
discretized real drilling cost	-0.021	-0.017	-0.025	-0.028	-0.036
	(0.014)	(0.016)	(0.026)	(0.030)	(0.028)
discretized real oil price	-0.007	-0.003	0.008	0.019	0.013
	(0.010)	(0.011)	(0.018)	(0.019)	(0.017)
mean development probability	0.043	0.041	0.039	0.036	0.029
# observations	1338	950	357	278	241

Notes: Marginal effects are reported. Standard errors in parentheses. Significance codes: * 5% level, ** 1% level, and *** 0.1% level.

FIGURE 5.

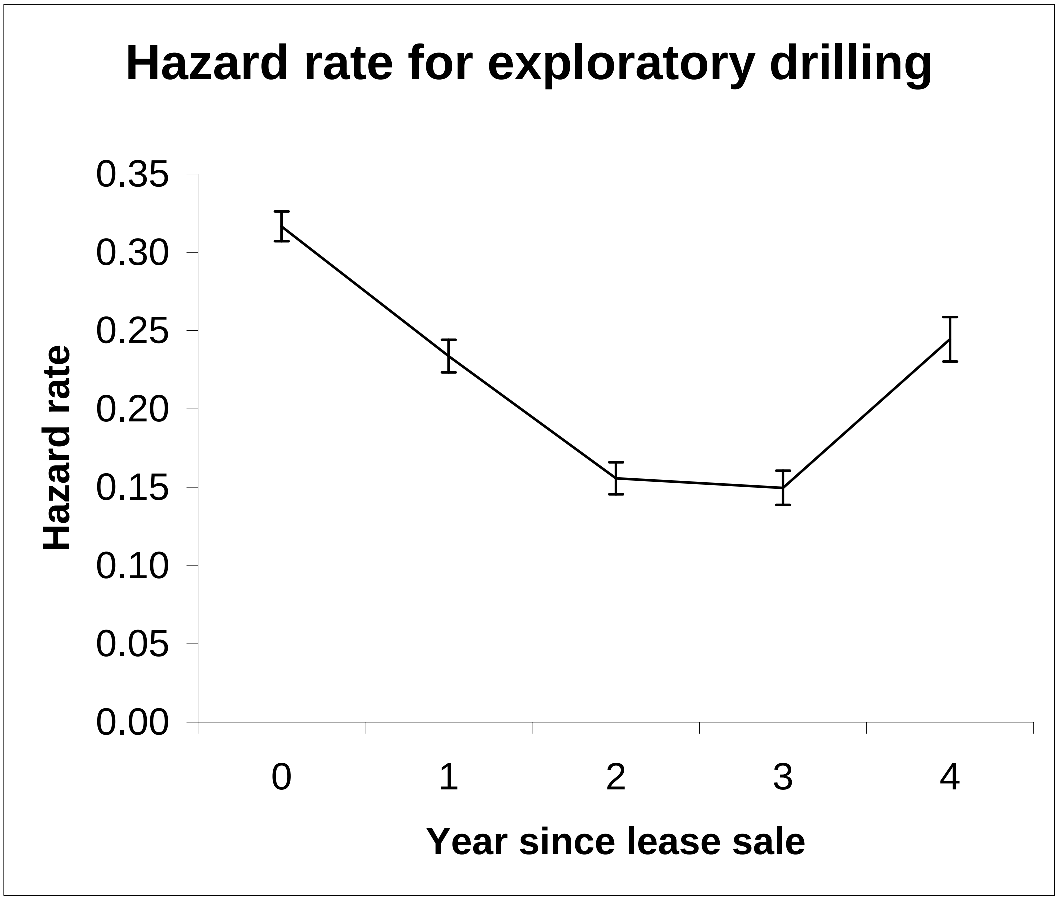

Notes: The sample consists of wildcat tracts sold before 1980 whose date of first drilling did not occur after the lease term. Error bars indicate plus or minus one standard deviation.
Source: Lin, 2009.

TABLE 6. Pooled results

	(1)	(2)
σ_μ	4.99	4.99
	(0.00)	(0.00)
σ_ε	4.86	4.96
	(0.03)	(0.03)
coefficient α in the exploration profit function on:		
discretized real drilling cost + 1	-10.01	-10.00
	(0.00)	(0.00)
coefficients γ in the development profit function on:		
other tract in market has been explored (dummy)	0.030	-0.068
	(0.023)	(0.051)
other tract in market has been developed (dummy)	0.16	0.16
	(0.02)	(0.03)
discretized average winning bid per acre	5.10	5.05
	(0.01)	(0.02)
discretized real drilling cost	-9.91	-9.97
	(0.01)	(0.03)
discretized real oil price	5.19	
	(0.00)	
real oil price in high bin		5.05
		(0.03)
real oil price in medium bin		5.02
		(0.01)
constant	5.01	4.96
	(0.02)	(0.03)
Moments		
exploration: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^e}(\Omega_{kt};\theta)-\overline{g^e(\Omega_{kt})}\right)\cdot n^{not\_yet\_e}(\Omega_{kt})$		
- model	0.00091	0.0010
- empirical	0.00038	0.00038
development: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^d}(\Omega_{kt};\theta)-\overline{g^d(\Omega_{kt})}\right)\cdot n^{e\_not\_yet\_d}(\Omega_{kt})$		
- model	0.0028	0.0032
- empirical	0.00074	0.00074
profits: $\frac{1}{n}\sum_{k,t}\left(E\left[\pi^d(\Omega_{kt},\varepsilon_{ikt};\theta)\mid I_{ikt}^d=1\right]-\overline{\pi^d(\Omega_{kt})}\right)\cdot n^d(\Omega_{kt})$		
- model	0.0074	0.0065
- empirical	-0.000084	-0.000084

Notes: Standard errors in parentheses. There are 1041 observations spanning 87 markets. Standard errors are formed by bootstrapping 100 simulated panels of 87 markets each. Additional moments are formed by interacting moments with state variables.

TABLE 7. Results for large tracts

	Acreage ≥ 5000
σ_μ	5.00
	(0.03)
σ_ε	4.90
	(2.67)
coefficient α in the exploration profit function on:	
discretized real drilling cost + 1	-10.00
	(0.01)
coefficients γ in the development profit function on:	
other tract in market has been explored (dummy)	0.0010
	(4.24)
other tract in market has been developed (dummy)	-0.0045
	(1.12)
discretized average winning bid per acre	5.07
	(1.28)
discretized real drilling cost	-9.97
	(0.80)
discretized real oil price	5.09
	(0.38)
constant	4.99
	(2.73)
Moments	
exploration: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^e}(\Omega_{kt};\theta)-\overline{g^e(\Omega_{kt})}\right)\cdot n^{not\_yet\_e}(\Omega_{kt})$	
- model	0.00052
- empirical	0.00041
development: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^d}(\Omega_{kt};\theta)-\overline{g^d(\Omega_{kt})}\right)\cdot n^{e\_not\_yet\_d}(\Omega_{kt})$	
- model	0.0038
- empirical	0.00086
profits: $\frac{1}{n}\sum_{k,t}\left(E\left[\pi^d(\Omega_{kt},\varepsilon_{ikt};\theta)\mid I^d_{ikt}=1\right]-\overline{\pi^d(\Omega_{kt})}\right)\cdot n^d(\Omega_{kt})$	
- model	0.0063
- empirical	-0.00013
# markets	55
# observations	698

Notes: Standard errors in parentheses. The acreage is the acreage of each tract in the market. Standard errors are formed by bootstrapping 100 simulated panels of size equal to the actual sample size. Additional moments are formed by interacting moments with state variables.

TABLE 8. Results for small tracts

	Acreage < 5000	< 4000	< 3000
σ_μ	5.73	5.00	3.48
	(0.03)	(0.13)	(0.00)
σ_ε	14.92	14.00	13.77
	(1.01)	(2.44)	(0.01)
coefficient α in the exploration profit function on:			
discretized real drilling cost + 1	-10.29	-10.01	-9.28
	(0.13)	(0.24)	(0.00)
coefficients γ in the development profit function on:			
other tract in market has been explored (dummy)	-25.78	-25.99	-27.67
	(0.15)	(0.56)	(0.00)
other tract in market has been developed (dummy)	3.15	4.00	4.06
	(0.03)	(0.17)	(0.00)
discretized average winning bid per acre	1.09	1.00	1.11
	(0.07)	(0.25)	(0.00)
discretized real drilling cost	-9.52	-10.00	-9.69
	(0.03)	(0.18)	(0.00)
discretized real oil price	14.65	16.00	15.93
	(0.07)	(0.34)	(0.00)
constant	-8.97	-9.00	-9.97
	(0.04)	(0.28)	(0.00)
Moments			
exploration: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^e}(\Omega_{kt};\theta)-\overline{g^e(\Omega_{kt})}\right)\cdot n^{not\_yet\_e}(\Omega_{kt})$			
- model	0.0016	0.0016	0.0020
- empirical	0.00070	0.00075	0.0010
development: $\frac{1}{n}\sum_{k,t}\left(\widehat{g^d}(\Omega_{kt};\theta)-\overline{g^d(\Omega_{kt})}\right)\cdot n^{e\_not\_yet\_d}(\Omega_{kt})$			
- model	0.00092	0.00094	0.0011
- empirical	0.0012	0.0017	0.0014
profits: $\frac{1}{n}\sum_{k,t}\left(E\left[\pi^d(\Omega_{kt},\varepsilon_{ikt};\theta)\,\vert\, I^d_{ikt}=1\right]-\overline{\pi^d(\Omega_{kt})}\right)\cdot n^d(\Omega_{kt})$			
- model	0.0105	0.0140	0.0013
- empirical	0.0000	0.0000	0.0000
# markets	29	21	19
# observations	308	247	210

Notes: Standard errors in parentheses. The acreage is the acreage of each tract in the market. Standard errors are formed by bootstrapping 100 simulated panels of size equal to the actual sample size. Additional moments are formed by interacting moments with state variables.

TABLE 9. Results of joint ownership simulation

	Acreage			
	All tracts	< 5000	< 4000	< 3000
number of years to exploration, conditional on exploring	0.038	2.47	2.47	2.45
	(0.14)	(0.22)	(0.34)	(0.26)
number of years to development, conditional on developing	12.51	2.47	2.47	2.45
	(1.93)	(0.23)	(0.34)	(0.26)
gross profits (million 1982 $), conditional on developing	14.48	9.73	8.63	7.43
	(0.07)	(1.07)	(2.69)	(0.27)
average additional gross profits per tract developed compared to separate ownership (million 1982 $)	25.31	17.00	17.28	16.51
# markets	87	29	21	19

Notes: The acreage is the acreage of each tract in the market. For each market, optimal policies are the joint exploration and development policies for both tracts that maximize joint expected profits over 100 simulated paths of the exogenous variables. Results are averaged over all markets for each of 100 draws from the estimated distribution of parameters. Standard deviations in parentheses.