

Hendricks, Kenneth; Pinkse, Joris; Porter, Robert H.

Working Paper

Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions

CSIO Working Paper, No. 0008

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Hendricks, Kenneth; Pinkse, Joris; Porter, Robert H. (1999) : Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions, CSIO Working Paper, No. 0008, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/38684>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0008

Empirical Implications of Equilibrium
Bidding in First-Price, Symmetric,
Common Value Auctions^{*}

By

Kenneth Hendricks, Joris Pinkse
Department of Economics, University of British Columbia

and

Robert H. Porter
Department of Economics, Northwestern University

August, 1999

^{*}We received helpful comments from Susan Athey, Phil Haile, Preston McAfee, and from the participants in a number of seminars. Hendricks and Pinkse received financial support SSHRCC, and Porter from the NSF.

Abstract

This paper studies federal auctions for wildcat leases on the Outer Continental Shelf from 1959 to 1970. These are leases where bidders privately acquire (at some cost) noisy, but equally informative signals about the amount of oil and gas that may be present. We develop a test of equilibrium bidding in a common values model that is implemented using data on bids and ex post values. We compute bid markups and rents under the alternative hypotheses of private and common values and find that the data are more consistent with the latter hypothesis. Finally, we use data on tract location and ex post values to test the comparative static prediction in common value auctions that bidders tend to bid less aggressively when they expect more competition.

1 Introduction

In previous work, [7],[8], we test equilibrium bidding in models with asymmetric information using drainage tract data. Drainage tracts are adjacent to tracts where oil and gas deposits have been found. The firms owning the adjacent tracts, called “neighbor” firms, have superior information to non-neighbors about tract value. The key feature of the theory of bidding between an informed firm and uninformed firms in a first-price sealed bid, common value auction is that the latter have to participate but their number is essentially irrelevant. Since neighbor firms could be identified, and because they behaved as a consortium, it was possible to test the theory by comparing the bidding behavior of neighbor and nonneighbors.

In this paper we study bidding in first-price, sealed bid auctions with symmetric information using wildcat tract data. Wells drilled in search of new deposits of oil and gas are called wildcat wells. Wildcat tracts are located in areas that have not been drilled. Firms are allowed to conduct seismic studies prior to bidding for these tracts, but they are not permitted to drill exploratory wells. The seismic studies yield noisy signals about the value of the tract. Thus, firms are more or less equally informed although they may have quite different opinions about the likelihood of finding oil and gas, depending upon the content and analyses of the surveys.

An important issue is whether the bidding environment is best characterized as one with common values or private values. The argument for the common value case is that the firms are uncertain about common components of the value of the lease being sold, such as the size of any oil or gas deposits under the tract, the prices of oil and gas over the likely production horizon if the lease is productive, and the common costs of exploration and development. The first component is likely to matter, because firms may have private information about the size of the deposits based on the seismic data they obtain, especially their interpretation of that data.

Alternatively, one might argue that there is little discrepancy in private assessments of these common components, and instead that valuations differ because of differences in bidder specific components of valuations. The most likely sources of bidder payoff heterogeneity are the private components of exploration and drilling costs. Bidders are not likely to differ in their valuation of recovered deposits, to the extent that there is a well developed market for oil and gas. Under this alternative view, valuations are best modelled as private, although they may be affiliated because of the common unknown components of payoffs that may be correlated with publicly available information.

The two different valuation models have quite different positive and normative implications. For example, an important strategic consideration is the level of competition, as measured by the number of potential bidders, say. In a private value model, bidding is more aggressive under more competition. In a common value envi-

ronment, this aggressiveness is tempered by winner’s curse considerations. Winning the auction is bad news to the extent that it reveals that the winning bidder’s signal of the lease’s value was more optimistic than that of the other bidders, and the greater the level of competition the worse the news associated with winning. One can construct examples where the winning bid is a decreasing function of the number of bidders. (See, for example, Klemperer [10].) The underlying model of valuations also matters in such normative issues as optimal auction design.

Despite the important theoretical distinction between common and private valuation models, the empirical literature has struggled with the problem of distinguishing between the two. There is a potentially serious identification problem. Laffont and Vuong [13] have shown that, conditional on the number of potential bidders, bidding data alone are insufficient to distinguish nonparametrically between a common value model and an affiliated private values model. Because bids in CV auctions are invariant to positive monotonic transformations of the signals, the common value environment is observationally equivalent to an affiliated private value environment. More precisely, suppose F is a joint distribution of signals and tract value that generates a bid distribution G when bidders bid according to the symmetric Bayesian Nash equilibrium. Then G can be rationalized by any distribution F' obtained from F by taking a positive monotonic transformation of the signals. Consequently, neither the bid function nor the underlying joint distribution function of signals can be identified in CV auctions in the absence of strong parametric restrictions on F . By contrast, both functions are identified (nonparametrically) in private value auctions.

One solution to the non-identification problem is to exploit the different comparative static results of the two models to distinguish between them. For example, Hong and Shum [9] assume a parametric form for the distribution of valuations under the two hypotheses and take advantage of the variation in the number of bidders to help identify the parameters of the distribution. Other papers have run reduced form regressions of bid levels on the number of bidders. They try to exploit the implication that individual bids in common value auctions should decrease as the number of participants in the auction gets relatively large, in contrast to the private values case. One potential problem with this approach arises when the number of bidders is employed as a regressor, instead of the number of potential bidders. In many auction settings, not all potential bidders choose to submit a bid, and the number of bidders should be treated as endogenous. In the Outer Continental Shelf (OCS) setting, it is the number of firms that search the tract and obtain signals that matters. If there is a binding reserve price, then not all participants necessarily bid. Consequently, the actual number of bidders may not be a good proxy for the level of competition. A second difficulty is finding variables that adequately capture the heterogeneity across tracts. In the OCS data, participation rates are typically less than 25 per cent in any given sale for any firm. Aggregate participation is positively correlated with ex

post tract value, as tracts with bigger deposits attract more bids on average. In the absence of good proxies for tract heterogeneity, ex ante bidder errors in expectation will be positively correlated with the number of bids submitted. That is, if bidders are more optimistic, more bid, and they bid aggressively.

Our contribution is two fold. First, we construct a measure of the set of potential bidders for individual wildcat tracts, and we examine how bidding varies with this measure. Our second, and more novel, contribution is to exploit information on ex post realizations of tract value. In particular, we construct a measure of ex post discounted revenues net of discounted royalty payments and drilling costs. We show how this information can be employed to distinguish between common and affiliated private values models and to test equilibrium bidding in the common value model. We argue that the OCS data are largely consistent with a common value model that accounts for a stochastic number of bidders, and inconsistent with a private values model. Furthermore, our test of equilibrium bidding in a CV environment is not rejected on the more competitive tracts.

The literature on structural estimation of auction models has focussed primarily on estimating the joint distribution of valuations of a fixed number of bidders using bid data. Smiley [21], Paarsch [19], and Donald and Paarsch [4] take a parametric approach, restricting the class of joint distributions to those which admit a closed form solution for the bid function, and then using maximum likelihood methods to estimate the unknown parameter vector. Baldwin, Marshall and Richard [3] and Marshall and Raiff [15] employ maximum likelihood methods to study collusion in an independent private values (IPV) environment. Bajari [1] uses Bayesian likelihood methods in an IPV model with asymmetric bidders. Laffont, Ossard and Vuong [12] develop a simulated non-linear least squares estimator which exploits the fact that, in the symmetric IPV environment, the bid function can be expressed as a conditional expectation of a second-order statistic. Elyakime, Laffont, Loisel and Vuong [5] employ a nonparametric approach to the estimation problem in an IPV environment and apply this method to timber sales. Li, Perrigne and Vuong [14] extend the nonparametric method to affiliated private values environments (APV) and apply their method to our data set. They assume that the dispersion in bids arises from idiosyncratic differences in costs. Bajari [2] develops a Bayesian econometric approach to structural estimation in an IPV environment based on an alternative solution concept, the quantile response equilibrium. His paper is the only paper on structural estimation other than ours that we are aware of that also considers the problem of censoring that arises when entry is not modeled. Our theoretical model is similar to the one developed by McAfee, Quan and Vincent [16] to study optimal reserve prices in real estate auctions.

The paper is organized as follows. In Section 2 we present the model. In Section 3 we describe how the federal government sells rights to oil and gas properties on

offshore lands and present the data. Section 4 provides evidence that bidders bid rationally. We classify tracts into two sets, a highly competitive set in which the number of potential bidders is larger than six and a less competitive set in which the number of potential bidders does not exceed six. For each category, we examine the relationship between actual returns and bids and provide a measure of the winner's curse. The results indicate that bidders bid less than their expectation of the value of the tract and less aggressively on the more competitive set of tracts. In Section 5 we develop and implement a test of equilibrium bidding in CV environment. In Section 6 we compare the implied rents and markups under the hypotheses of private and common values models. Concluding remarks are contained in Section 7.

2 The Model

Consider a first-price sealed bid auction in which l potential bidders compete for the production rights of a single tract. The announced minimum bid is r . Prior to bidding, each potential bidder makes a private, independent decision whether to invest in a seismic survey. Firms that invest in such surveys are called *active*. We adopt the convention of denoting random variables in upper case and realizations of random variables in lower case. Let N denote the number of active bidders. We assume that the probability of being active is the same for each potential bidder. Hence, N is a binomial random variable with parameters (l, λ) , where λ represents the probability that a potential bidder is active. Firms that are not active do not bid.

The unknown common value of the oil and gas deposit is denoted by V . Each active bidder i receives a real-valued information signal s_i that affects his valuation. The actual value of the tract to bidder i is given by $U_i = u(V, S_i)$ where u is non-negative, continuous and increasing in both arguments. All of the bidder utilities depend upon the common component in the same manner, and each bidder's utility is allowed to depend upon its own signal. The random variables (V, S_1, \dots, S_l) are (positively) affiliated which, loosely speaking, implies that higher realizations of some of the components make higher realizations of the other components more likely. This condition is essential for the existence of a Bayesian Nash equilibrium in strictly increasing bid functions. We shall also assume that, conditional on V , the bidders' signals are independent.¹ Let F denote the cumulative probability function of the $l + 1$ vector (v, s_1, \dots, s_l) . It is assumed to be symmetric in the information signals and to have a density $f(v, s_1, \dots, s_l)$. If $n < l$, then the joint distribution of the $n + 1$ vector (v, s_1, \dots, s_n) is derived from F by integrating over (S_{n+1}, \dots, S_l) . Finally, we assume

¹McAfee, Quan and Vincent [16] note that conditional independence is actually a consequence of the exchangeability of an unbounded number of random variables and thus, for n large, is a reasonable restriction on the distribution.

that N is independent of the random variables (V, S_1, \dots, S_l) .

The probability law of N , the distribution function F and the utility function u are common knowledge. Each active bidder i knows the value of his signal s_i but does not know the signals of the other active bidders. Anecdotal evidence suggests that the decision to acquire detailed survey information is private,² in which case the number of active bidders is not known to the bidders. For the purposes of the theoretical model, we shall assume that bidders do not observe n . However, our empirical model is consistent with either informational assumption.

Our model is similar to that of McAfee, Quan and Vincent [16], who extend the model referred to by Laffont and Vuong [13] as the Affiliated Value (AV) model to the case of a stochastic number of bidders. The AV model was first introduced by Wilson [23] and is a special case of the general symmetric model of Milgrom and Weber [18]. In the AV model, the signals of other bidders affect the expected utility of bidder i through their affiliation with v , but they do not enter as arguments of the utility function. Two special cases will be of interest. The model is said to be a Common Value (CV) model when $U_i = V$; it is called an Affiliated Private Value (APV) model when $U_i = S_i$.

In what follows, we restrict attention to strictly increasing differentiable symmetric Bayesian Nash equilibria. As McAfee, Quan and Vincent [16] have observed, when the number of bidders is stochastic, the existence of such equilibria is not guaranteed without further restrictions on the joint distribution F . They show that the following condition on F is sufficient, namely,

$$[1 - \lambda(1 - F_{S|V}(s|v))]/f_{S|V}(s|v)$$

is decreasing in v . Note that, when $\lambda = 1$, the condition is implied by affiliation of (S, V) . However, for $\lambda < 1$, it is unlikely to be satisfied near the lower bound of the support of S . McAfee and Vincent argue that a binding reserve price makes the condition more plausible, and we shall assume that this is the case for the oil lease auctions.

In the derivation below, we take the perspective of bidder 1. But before specifying his optimization problem, we require some notation. Let $K = N - 1$ denote the number of rival bidders and let

$$p_k = \binom{l-1}{k} \lambda^k (1-\lambda)^{l-1-k}$$

denote the probability that bidder 1 faces k rival bidders. Define Y_1 as the maximum signal among bidder 1's rivals conditional on the event that bidder 1 has at least

²For example, some of the potential bidders may purchase geological survey information from corporations that specialize in such activities. Other bidders, such as Shell, typically conduct their own surveys.

one rival. Let $F_{Y_1|S_1}(\cdot|s)$ denote the cumulative distribution of Y_1 when bidder 1 has obtained signal s and let $f_{Y_1|S_1}(\cdot|s)$ denote the associated probability density function. Here,

$$F_{Y_1|S_1}(y|S_1 = s) = \sum_{k=1}^{l-1} \frac{p_k}{1 - p_0} F_{Y_1|S_1,K}(y|S_1 = s, K = k)$$

Note that the probability weights have been normalized to sum to 1 by conditioning on the event that bidder 1 has at least one rival. For each k , affiliation implies that $F_{Y_1|S_1,K}$ is nonincreasing in s . Hence, $F_{Y_1|S_1}$ is also nonincreasing in s . Finally, define

$$w(s, y) = E[u(V, s)|S_1 = s, Y_1 = y]$$

to be bidder 1's expected utility from owning the tract when his signal is s and the maximum of his rivals' signals (assuming $k \geq 1$) is y . When bidder 1 has no rivals, his expected utility is given by

$$w(s) = E[u(V, s)|S_1 = s].$$

Now suppose that each rival adopts the monotone increasing bidding strategy $\beta(s)$ with inverse $\eta(b)$. Under the assumption of risk neutrality, bidder 1's optimization problem consists of choosing b to maximize

$$\Pi(b, s) = (1 - p_0) \int^{\eta(b)} (w(s, y) - b) f_{Y_1|S_1}(y|s) dy + p_0(w(s) - b).$$

The first-order condition for a maximum is

$$(1 - p_0)[(w(s, \eta(b)) - b) f_{Y_1|S_1}(\eta(b)|s) \eta'(b) - F_{Y_1|S_1}(\eta(b)|s)] - p_0 = 0. \quad (1)$$

If bidder 1's best reply is $b = \beta(s)$, then, substituting for b , equation (1) can be expressed as a differential equation:

$$(1 - p_0)[(w(s, s) - \beta(s)) \frac{f_{Y_1|S_1}(s|s)}{\beta'(s)} - F_{Y_1|S_1}(s|s)] - p_0 = 0. \quad (2)$$

This differential equation can be solved subject to a boundary condition. Define

$$s^*(r) = \inf \{s | (1 - p_0) E[w(s, Y_1)|S_1 = s, Y_1 < s] + p_0 w(s) \geq r\}$$

to be the lowest signal at which a bidder believes the value of the tract conditional on winning (in a symmetric equilibrium) is worth at least the reserve price. We assume that $s^*(r)$ exceeds the lower bound of the support of S . Hence, the reserve price

is binding and screens some buyers from participating. In that case, the boundary condition for the differential equation given above is $\beta(s^*) = r$.

A parametric approach would be to assume risk neutrality, as we have above, specify $F_{Y_1|S_1}$ in terms of the primitives λ, F , and then estimate λ and the unknown parameters of F by maximizing a likelihood function of observed bids based on the equilibrium relation. The approach is complicated (at least for us) since it requires solving for the inverse bid function, which typically does not have a closed form solution and has to be estimated numerically. In addition, as we shall see later, there may be an identification problem.

Laffont and Vuong's important insight is that for empirical purposes it is more useful to invert the above equilibrium relation and express the signal as a function of the bid and the distribution of bids. Define

$$M_1 = \max\{r, \beta(Y_1)\}$$

to be the highest bid submitted by bidder 1's rivals or, if no rivals submit a bid, the reserve price. Let the conditional distribution of M_1 given B_1 be denoted by $G_{M_1|B_1}(\cdot|\cdot)$ and its density by $g_{M_1|B_1}(\cdot|\cdot)$. Note that monotonicity of β, η implies, for any $b > r$,

$$G_{M_1|B_1}(m|b) = (1 - p_0)F_{Y_1|S_1}(\eta(m)|\eta(b)) + p_0$$

Here $G_{M_1|B_1}(m|b)$ is the probability that the highest bid among bidder 1's rivals is less than m conditional upon bidder 1's bid of b . The associated density function is given by

$$g_{M_1|B_1}(m|b) = \frac{(1 - p_0)f_{Y_1|S_1}(\eta(m)|\eta(b))}{\beta'(\eta(m))}.$$

Substituting the above relations into equation (2) yields

$$w(s, s) = b + \frac{G_{M_1|B_1}(b|b)}{g_{M_1|B_1}(b|b)} \equiv \xi(b, G). \quad (3)$$

Note that the expression on the right-hand side of equation (3) depends only upon bids and can be estimated directly from bid data. Affiliation implies that $w(s, s)$ is a monotone increasing function of s . However, its interpretation depends critically upon assumptions about preferences. If valuations are assumed to be private, then

$$w(s, s) = s.$$

Alternatively, if valuations are common, then

$$w(s, s) = E[V|S_1 = s, Y_1 = s].$$

One of the key challenges to empirical work in auctions is to distinguish between the two interpretations. We will return to this issue in a later section.

3 Auction Mechanism and Data

The U.S. government holds the mineral rights to offshore lands more than three miles from the coast, out to the 200 mile limit. The states own the rights out to the three mile mark. Beginning in 1954, the federal government has transferred production rights on its lands to the private sector by a succession of lease sales in which hundreds of leases have been auctioned. A wildcat lease sale is initiated when the U.S. Department of Interior (DOI) announces that certain offshore areas are available for exploration, and nominations are invited as to which tracts should be offered for sale. A tract is typically a block of 5,000 or 5,760 acres, or half a block. The number of tracts available in a sale is usually well over one hundred and tracts are often scattered over several different areas.

Prior to the sale, and often prior to the announcement of the sale, firms conduct seismic surveys on selected tracts. The time between announcement and sale date is at most several months, since DOI is required to give only thirty days notice. A survey provides information about the geology of the tracts and is used by a firm to determine which tracts to bid and how much to bid. The cost of a detailed survey has been reported to vary between \$9 to \$26 per acre, including payments to the geologists to study and interpret the seismic data. Some firms conduct their own surveys while others purchase surveys from geophysical firms that specialize in this type of activity. Many aspects of the firm's tract evaluations are private, including which tracts were selected for investigation. In particular, given the cost of a tract survey, it is not an equilibrium for all firms known to be interested in acquiring leases in an area to survey all tracts. Individual firms typically select a fraction of the available tracts and the locations of these tracts are reputedly a closely guarded secret.

A firm can choose either to bid solo or jointly with other firms. A joint bid in offshore auctions consists of two or more firms combining to submit a single bid and sharing the costs and revenues if their bid is the high bid. Prior to 1975, this practice was legal for all firms. In late 1975, DOI adopted regulations barring the eight largest crude oil producers worldwide (Exxon, Gulf, Mobil, Shell, Standard Oil of California, Standard Oil of Indiana, Texaco and British Petroleum) from bidding with each other on the grounds that the practice was reducing prices. Joint bidding groups are frequently sale-specific and form after the announcement of the sale.

The nominated tracts in a sale are sold simultaneously in a first-price, sealed bid auction. The announced reserve price for tracts in our sample is \$15 per acre. A participating bidder or consortium of bidders submits a separate bid on each tract that it has an interest in acquiring. A bid is a dollar figure, known as a bonus. At the sale date, DOI opens the envelopes and announces the value of the bids that have been submitted on each tract and the identities of the bidders. The firm or consortium that submits the highest bid on a tract is usually awarded the tract at a

price equal to its bid. In practice, the government could and did reject bids above the stated minimum price. The rejection rate was less than 10% on wildcat tracts and usually occurred on marginal tracts receiving only one bid (Porter [20]). We will largely ignore this factor in our analysis.

When a firm or consortium is awarded a tract, it has 5 years to explore it. If no work is done during the lease term (i.e., no wells are drilled), ownership reverts to the government and the tract may subsequently be re-offered. A nominal fee, typically \$3 per acre, is paid each year until either the lease is relinquished or production begins. If oil or gas is discovered in sufficient quantities, the lease is automatically renewed as long as production occurs. A fixed fraction of the revenues from extraction accrues to the government as royalty payments. The royalty rate for tracts in our sample is 1/6.

We restrict attention to sales of wildcat tracts off the coasts of Texas and Louisiana held during the period 1954 to 1970 inclusive. The information available for each tract receiving at least one bid are the date of sale; acreage; location; the identity of all bidders and the amounts they bid; the identity of participants in joint bids and their shares in the bid; whether the government accepted the high bid; the number and date of any wells that were drilled; and monthly production through 1991 of oil, condensate, natural gas, and other hydrocarbons. The annual survey of drilling costs conducted by the American Petroleum Institute is used to compute drilling costs. Production flows are converted into revenues using the real wellhead prices at the date of the sale, and discounted to the auction date at a 5 percent per annum rate. The ex post value of a tract is defined to be discounted revenues less discounted drilling costs and royalty payments.

3.1 Number of Potential Bidders

Our measure of l , the number of potential bidders on a tract, is constructed from information on who bid in the area and when. For tracts that were drilled, location is identified by the longitudinal and latitudinal coordinates of the well. Tracts that were not drilled are assigned coordinates by interpolation from nearby tracts that were drilled.³ On average a tract covers 0.0463 degrees of longitude and 0.0405 degrees of latitude. A *neighborhood for tract t* consists of all tracts whose registered locations are within 0.1158 (2.5 times 0.0463) degrees of longitude and 0.1012 (2.5 times 0.0405) degrees of latitude of tract t and that were bid on at the same time or prior to tract t . We denote the neighborhood by Λ_t . Ignoring irregular tract sizes and boundary effects, the maximum possible size of Λ_t is 25 tracts or 125,000 acres.

³A small number of tracts were sufficiently isolated that it was not possible to interpolate their location from nearby tracts. These tracts were dropped from the sample.

An obvious approach to defining the number of potential bidders on a tract is simply to count the number of firms that bid on the tract or in its neighborhood. The rationale is that if a firm is interested in the area, then it will probably bid on at least one tract. In practice, two difficulties arise with this measure. First, many firms bid infrequently, and they were probably not perceived as serious competitors by the major bidders. Second, even if we restrict attention to firms that bid frequently, they often bid jointly.

Table 1 lists the bidding activities of the twelve firms and bidding consortia with the highest participation rates in our sample. Three bidding consortia pooled their exploration budgets and expertise and bid almost exclusively with each other in each sale. We treat these consortia as single firms. The twelve firms and consortia listed in the table are designated as *large* firms. All other firms are referred to as *fringe* firms. For the purposes of this paper, we define a joint bid as one in which two or more large firms participated. All other bids are called *solo* bids. The first two columns of Table 1 give the number of solo and joint bids of each large firm.⁴ The twelve large firms account for about 75% of all bids in our sample.

We construct our measure of the number of potential bidders on a tract t as follows. First, we restrict the set of potential bidders to the set of large firms that bid, either solo or jointly, on at least one tract in Λ_t (which includes tract t). Let Γ_t denote this set of firms. The cardinality of this set serves as an upper bound for our measure of competition on tract t . Firms that submit solo bids on tract t are counted as competitors since they are revealed to be active. Firms that submitted a joint bid on tract t are treated as a single competitor, regardless of how they bid on other tracts in Λ_t , under the assumption that their joint venture is known to competitors. We also include firms that did not bid on tract t but submitted solo bids on at least one tract in Λ_t . Firms that did not bid on tract t and submitted only joint bids with each other on tracts in Λ_t are also counted as single competitors.⁵

The number of potential bidders obtained from applying the above criteria to firms in Γ_t may overstate the true level of competition. Solo bids are always treated as evidence of competitive behavior. But firms could coordinate bidding strategies by agreeing to bid solo on different sets of tracts rather than bidding jointly. Our measure does not capture this form of collusion. On the other hand, firms known to be interested in the area but who decided not to bid on any tracts in Λ_t are not

⁴A small number of tracts registered multiple bids by a firm. This problem arose in part due to classification errors in identifying a firm's subsidiaries and affiliates. We adopted the following rule for these bids. If a subset of the participants in one bid participated in another bid, the latter is dropped. Thus, solo bids of bidders who also submitted joint bids are eliminated. In the other cases, the highest bid is taken and the others dropped.

⁵We did calculate the number of potential bidders under the alternative hypothesis that joint ventures are private. However, the qualitative results did not change.

counted. We might then underestimate the number of potential bidders.

The third column of numbers in Table 1 reports the number of tracts where each large firm is counted as a potential bidder. The fourth column gives each firm's bid participation rate on tracts where it is a potential bidder. Note that the variation in these participation rates across firms is considerably less than the variation in their bid frequency rates (i.e., the number in column 3 divided by 1260, the number of tracts in the sample). The participation rates of most of the firms falls between 30 and 47 per cent. The AGCC consortium is the outlier with a participation rate of 54%. In contrast, bid frequency rates vary more or less uniformly between 13 per cent (Phillips) and 44 per cent (AGCC).

3.2 Sample Statistics

Table 2 provides summary statistics on the sales in our sample. Typically, at least one bid was submitted on approximately 50 per cent of the tracts offered in a sale. Big12 firms typically bid on over 80 per cent of the tracts in any given sale. The government rejected the high bid on 7 per cent of the tracts receiving bids. The unsold tracts are mostly in later sales, and are mostly marginal tracts. The fraction of tracts drilled and the fraction of hits (i.e., productive tracts) among those that were sold do not vary much across the larger sales and are typically about 75% and 50% respectively. Mean discounted revenues on productive tracts are similar across the larger sales, with the notable exception of the sale in 1968 where revenues were only \$12.21 million per tract. (All dollar magnitudes are in 1982 dollars.) Mean discounted net revenues are net of royalty payments and discounted costs. They are calculated for all tracts, including "dry" tracts. Mean net revenues vary considerably across the sales, averaging \$10.5 million per tract across the sales. The average winning bid for the entire sample is \$6.2 million, which yields an average "rent" of \$4.3 million per tract. The average winning bid was substantially higher in later sales.

Table 3 provides summary statistics on wildcat tracts. The tracts are classified by the number of potential bidders, which ranges from 0 to 12. For each value of l , the first row gives the number of tracts, the second gives the number of bids per tract, and the third gives the number of Big12 bids per tract. The frequency distribution is approximately bi-modal, with peaks at $l = 3$ and $l = 9$, and a median of 7. The number of bids is positively correlated with, but often considerably smaller than, the number of potential bidders. Even when all of the Big12 bidders are potential bidders, the average number of bids is only 4.25. Note the sharp increase in the average number of bids at the median, from 2.80 to 3.94. A comparison of rows two and three reveals that the probability of a fringe bid is increasing with l , and on average there is one fringe bid. This suggests that excluding fringe firms is probably not introducing too much error in our measure of competition. Ex ante expectations,

as measured by the high bid, are positively correlated with l . The average high bid increases from \$600 thousand on tracts where none of the Big12 firms are potential bidders to \$13.1 million on tracts where every Big12 firm is a potential bidder. Note that there is a relatively large increase in the level of the high bid when l increases from 6 to 7. The percentage of tracts drilled increases with the number of potential bidders but the hit rate, defined relative to the number of tracts drilled, appears to be independent of the number of potential bidders for l greater than 2, fluctuating between 40 and 55%. Average revenue on productive tracts is quite noisy and does not exhibit any trend as l increases. Net revenue is higher on the more competitive tracts but once again the relationship between these variables and l is quite noisy and not monotonic. By contrast, average hit rates and revenues are much more strongly correlated with the number of bids (Porter [20]).

We also computed the average high bid, acreage bid, hit rate, average revenue and net revenue for neighborhoods of tracts with l potential bidders. The neighborhood variables all tend to increase with l , which reflects the spatial correlation in deposits. In any region, one or two tracts receive most of the action and competition tends to diminish the further away one gets from the center of interest. The fraction of acreage sold as wildcat tracts before or after tract t is typically quite small. One reason is that the federal offshore lands have been explored in a series of bands that extend along the coastline and move outward from the shoreline into the Gulf of Mexico. As a result, most of the tracts in Λ_t that were sold as wildcat tracts were sold in the same sale as tract t . The other reason is that any tract in Λ_t sold in a later sale is likely to be classified as a drainage or development tract.

In the analysis that follows, we stratify the sample according to the number of potential bidders into two categories, high and low. The low category is defined by $l \leq 6$, and the high category by $l > 6$. The highly competitive set has 752 tracts and the less competitive set has 508 tracts. Table 4 presents the frequency distribution of (base 10) log bid for each category. Not surprisingly, the low l distribution has many fewer bids than the high l distribution. Most of the density of the low l distribution lies in the range 5.4 to 7.0, which corresponds to bids of \$250 thousand to \$10 million respectively. The main difference between the low l and high l distributions, apart from the number of bids, is that the latter contains relatively more high bids. There is a substantial number of bids in the interval 7.0 to 7.8 (\$10 million to \$60 million) on high l tracts.

The classification of tracts into highly competitive and less competitive sets accounts for some tract heterogeneity. The following table, which is related to Table 3,

illustrates this point.

	Hibid	Drill Rate	Hit Rate	NetRev
Low l	2.76	70.3	45.6	8.00
High l	8.51	81.7	48.4	12.12

Clearly, the high l tracts are more likely to be productive than low l tracts. Net revenues per productive tract are 52% higher on the more competitive set than on the less competitive set. This partially explains why the average high bid is 84% higher on high l tracts than on low l tracts.

4 Rationality and the Winner’s Curse

In this section we conduct some basic tests of bidder rationality, and construct a measure of the magnitude of the winner’s curse as a function of the level of competition. We also examine the data for evidence that bidders bid less aggressively in auctions with more potential bidders. Throughout this section and the remainder of the paper, we will drop the subscript and let B and M denote, respectively, the bid of the representative bidder and the maximum bid of its rivals.

The weakest possible behavioral test of rational bidding is that bidders should bid less than their expected tract value. Let Z denote an index of observable (to the bidders) tract characteristics. Under the hypothesis of a common value environment, we can express this restriction as

$$E[V|B = b, Z = z] > b. \tag{T1}$$

We refer to this test as (T1). Rational bidders in a common value environment should also anticipate the “bad news” associated with winning. That is, they should bid less than their expected value of the tract conditional on winning,

$$E[V|B = b, M < b, Z = z] > b. \tag{T2}$$

We refer to this test as (T2). The difference between these two conditional expectations is a measure of the “winner’s curse”,

$$\kappa(b) = E[V|B = b, Z = z] - E[V|B = b, M < b, Z = z].$$

The conditional expectations can be estimated from data on bids and net revenue, assuming the latter is a good proxy for the common component in bidder valuations. That is, observed tract revenues and costs are assumed to be the same regardless of who actually won the tract. Note that both of the above tests and our measure of

“winner’s curse” can be applied to individual bidders, or to a set of bidders. Furthermore, in a symmetric common values environment, the winner’s curse measure should be greater when there is more competition, as winning is worse news the larger the number of potential bidders.

Figure 1 plots estimates of the average value of V conditional on b and the average value conditional on the event of winning at b for low l tracts, where there is less potential competition. In both cases, the estimates of the conditional expectations are computed from a univariate locally linear regression. The former estimate, which is the basis for (T1), employs the full sample of bids submitted by Big12 firms, whereas the latter uses the subsample of winning bids by the Big12. Details of the nonparametric estimation procedures can be found in the Appendix. We also plot a 45 degree line, so that these conditional expectations can be compared to the relevant bid level. We use a (base 10) log scale for both axes of the figure. Both conditional expectation functions are above the 45 degree line over the observed range of bids, which means that Big12 bidders satisfy the rationality tests (T1) and (T2) in the aggregate. Throughout the range of bids, the average value of tracts won by large firms at b is substantially lower than the average value of tracts that received a bid of b . Furthermore, the difference increases weakly with b . The difference is depicted in levels in Figure 1a, as a function of log bid. For example, at a bid of \$1 million, the difference is on the order of \$2 million. A block bootstrap procedure, described in Appendix B, is employed to derive a one-sided 95% confidence band for our estimate of the winner’s curse. The confidence band depicted in Figure 1a indicates that the difference between the two expectation functions is marginally significant at the 5% level.

Figure 2 plots estimates of the same two conditional expectations for high l tracts. The results are qualitatively similar to those reported for low l tracts. Figure 2a presents our estimate of “winner’s curse” and its confidence band for tracts with more competition. Measured relative to bid, the “winner’s curse” appears to be larger for the high l tracts than the low l tracts. For example, the difference between bidding \$1 million and winning at \$1 million is on the order of \$5 million on high l tracts. Figure 2a indicates that the difference between the two expectation functions is significant for bids less than \$1 million, but marginally significant at higher bid levels. It suggests that the “winner’s curse” is present and increases weakly with bid. However, most of the large firms appear to anticipate the “curse” and bid on the order of 1/3rd of the average value conditional on winning.

In summary, the analysis of aggregate bidding on the part of Big12 firms indicates that they pass our tests of rational bidding, (T1) and (T2).⁶ Averaged across all bids,

⁶We also computed the conditional expectations for individual firms and found that most firms bid rationally. However, Texaco not only failed to anticipate the “winner’s curse”, it also failed to

we estimate the value of the winner’s curse to be \$2.73 million on low l tracts, and \$6.13 million on high l tracts. The respective standard errors are \$1.18 million and \$1.61 million. The estimate of winner’s curse is 106.5% of the average winning bid on low l tracts, and 75.3% on high l tracts.

One may be tempted to conclude from these results that the bidding environment is common value. Absent unobserved heterogeneity, this conclusion would be correct since winning should not affect valuations in a private values environment. But, to see why unobserved heterogeneity matters, consider an APV setting. At a given bid level b , the bid is more likely to win if the idiosyncratic component of valuation is favorable. Hence the expected common component of valuations is lower conditional on winning. That is, there is a selection effect associated with conditioning on winning. Furthermore, the magnitude of the selection effect is higher on high l tracts than on low l tracts.

4.1 Comparative Statics

We now examine the prediction that bidders bid less aggressively on tracts with more bidders. As mentioned in the introduction, the key empirical issues in testing this prediction are (i) measuring the number of bidders and (ii) controlling for tract heterogeneity. The number of bids is the correct measure of competition if it is common knowledge which firms are interested in bidding on individual tracts, and if all interested firms in fact submit bids. However, in what follows, we shall use the number of potential bidders as our primary measure of competition since we suspect that the number of active bidders is not observed and, even if it is observed, it is likely to be correlated with tract characteristics that are not observed by the econometrician. The unobserved tract heterogeneity remains the main difficulty. This became obvious to us when we computed the empirical bid distributions conditional on l and/or n . Bidders bid more on average on tracts with higher values of l and/or n . Furthermore, this behavior is not unreasonable since the empirical distribution of V is also stochastically increasing in l and n .

Our approach to resolving the tract heterogeneity problem is to impose a restriction on the joint distribution of (S, V) that allows us to identify the bid function. Wilson [23] adopts the normalization $E[S|V = v] = v$, where signals are measured so that the mean signal on a tract is equal to the tract’s value. We instead assume that

$$E[V|S = s, Z = z] = s. \tag{R}$$

Condition (R) states that if a firm obtains a signal s on a tract with characteristics z , then the expected value of that tract is equal to the value of the signal. Signals are bid less than the average value conditional on its bid. Its $E[V|b]$ curve lies everywhere below b .

normalized in terms of ex post value and we assume that firms' posterior estimates are not biased but correct on average. Identification of the bid function follows immediately from condition (R) and monotonicity of the bid function since

$$E[V|S = s, Z = z] = E[V|B = \beta(s, z), Z = z] \implies \eta(b, z) = E[V|B = b, Z = z] \quad (4)$$

Equation (4) defines an inverse bid function. It can be estimated as follows. For every bid level b on a tract with characteristics z , define a neighborhood $B(z)$ of b , and compute the average ex post value of all tracts with characteristics z that received a bid in $B(z)$. To implement this idea, we employ a kernel estimator of the mean ex post value in the neighborhood of any bid b for tracts with similar characteristics.

How does this approach finesse the problem of unobserved tract heterogeneity? Condition (R) is more than a normalization assumption on the distribution of signals. It also implies that S is a sufficient statistic for V in the conditional distribution of V given S and Z . Thus, tract heterogeneity may influence the participation decision of potential bidders and, if surveys are private, the beliefs of an active bidder about the number of rivals, but it does not affect an active bidder's beliefs about V . (Note that the assumption also allows us to ignore what may otherwise be a serious sample selection problem.) If the sample is stratified according to some characteristic z_1 , then the condition is a statement about the joint distribution of V, S , and other characteristics that are observed by the bidders, but not necessarily by the econometrician, conditional on z_1 .

Figure 3 presents the estimates of the bid functions for tracts with large and small numbers of potential bidders. That is, in Figure 3 we restrict our attention to one observable tract characteristic z , whether or not the number of potential bidders is relatively large, and stratify the sample according to that characteristic. The results indicate that firms bid somewhat less aggressively on high l tracts than low l tracts for a given signal, consistent with a common values environment, but inconsistent with a private values setting. Figure 4 presents the corresponding estimate of the probability density function of private signals for high and low l tracts. Clearly, the distribution of signals on high l tracts stochastically dominates (in the first order sense) the distribution of signals on the low l tracts.

In some sense, our identifying assumption leads us to a non-parametric reverse regression of ex post value on bid and tract characteristics. The alternative identifying restriction, that the conditional mean of signals is the ex post value (i.e., $E[S|V = v] = v$, or $E[S|V = v, Z = z] = v$) would lead us to regress bid on value and tract characteristics. One problem with these alternative conditions are that they are unlikely to be satisfied for high values of V (as well as zero values). In our data, there is less variability in bids than in our measure of ex post value, which equals zero on many tracts. A second problem with this alternative method is that ex post value is measured with error. One would favor the results of the reverse regression under

the supposition that the measurement error is relatively severe, and, consistent with this supposition, the estimated slope of the bid function from the regression of bid on value is much less than that from the reverse regression.

These comparative statics results are suggestive, rather than definitive, because of the maintained assumption about the sufficiency of the signal, condition (R). We also computed the relationship between bids and ex post payoffs conditional on the local neighborhood variables described in Section 3.2, in addition to conditioning on our measure of competition, but these results were not informative. A problem with the neighborhood variables is that they combine both ex ante and ex post information. The latter characteristic is problematic for the purpose of computing bid functions, given the spatial correlation in information and returns. Nevertheless, these neighborhood variables may be useful in constructing bid participation models.

5 Tests of Equilibrium Bidding

In this section we describe two tests of equilibrium bidding based on equation (3) and report the results of these tests for the OCS data. The maintained hypothesis throughout this section is that the bidding environment is common values.

Recall that under the assumption that signals are affiliated, $w(s, s)$ is increasing in s . Hence, as Laffont and Vuong [13] have observed, one implication of equilibrium behavior is monotonicity of ξ . If $\xi(b, G)$ is not monotone increasing in b , then bidding behavior is not consistent with a symmetric Bayesian Nash equilibrium. We test this monotonicity restriction below.⁷

Our second test is based upon the assumption that our measure of discounted tract net profits is a good proxy for the common value V . Define

$$\zeta(b, z) = E[V|B = b, M = b, Z = z].$$

It then follows from monotonicity of β and equation (3) that if the bidders bid according to a Bayesian Nash equilibrium,

$$\zeta(b, z) = \xi(b, G(z)). \tag{5}$$

An important feature of this test is that it is not sensitive to unobserved tract heterogeneity. Since equality holds at every realization of Z , it must also hold in the aggregate. We implement the test separately for high l and low l tracts.

For each category of l , the bid data used to compute the conditional expectation ζ and the distribution $G_{M|B}$ are constructed as follows. Let b_{it} denote the bid submitted

⁷Note that this test does not require any commitment on whether the auction environment is private or common values.

by bidder i on tract t and let m_{it} denote the maximum bid submitted on tract t by bidders other than i . If no other bid was submitted, then m_{it} is set equal to the reserve price. We select only pairs (b_{it}, m_{it}) where i indexes one of the Big12 firms. However, *all* bidders, including fringe bidders, are used in determining m_{it} . The estimates of the conditional expectation ζ are computed from a bivariate locally linear regression. The estimates for ξ are obtained from an estimate of the distribution $G_{M|B}$.

Figure 5 depicts the estimates $\hat{\xi}$ and $\hat{\zeta}$ for low l tracts. We use a (base 10) log scale for both axes of the Figure. We also plot a 45 degree line, so that $\hat{\xi}$ and $\hat{\zeta}$ can be compared to the relevant bid level. At any bid level b , the vertical difference between $\hat{\xi}(b)$ and the 45 degree line represents the factor by which bidders' markdown their bid from their conditional expectation of tract value. The difference should be positive and Figure 5 reveals that this is the case throughout the range of bids. Note that the markdown in this case is relative to the bidder's expected value conditional on the event that the maximum signal among its rivals is *equal* to its own signal, not conditional on the event of winning. It should not be interpreted as a measure of the bidder's expected profit. The second point to note about $\hat{\xi}$ is that it is strictly increasing throughout the range of bids. Hence, the above model passes our first test of equilibrium bidding, at least for the set of tracts where competition is low. The conditional expectation $\hat{\zeta}$ also lies above the 45 degree line for the relevant range of bids and is increasing throughout. However, the difference $\hat{\zeta} - \hat{\xi}$ does not appear to be close to zero, except for bids near \$1 million. At higher bid levels, bidders appear to overestimate the value of the tract at the time of bidding.

Clearly, a formal test of equality of $\hat{\xi}$ and $\hat{\zeta}$ is needed. Such a test would probably contain elements of tests for the hypothesis that a nonparametric regression function is zero everywhere. It requires the derivation of the asymptotic distribution of an estimate of the (properly renormed) difference $\zeta - \xi$. In addition, since the small sample properties of such a test are likely to be poor a bootstrap correction may need to be applied. There is a fairly large literature on bootstrapped confidence bands of nonparametric regression estimators (see e.g. Hall [6] and the references therein). Our estimator does not fit any of the standard cases and obtaining *asymptotic refinements*, i.e. confidence bands more accurate than those obtained by first order asymptotics, is very difficult. We are satisfied with obtaining confidence bands which are asymptotically valid. To account for spatial dependence, we employ a new block bootstrap procedure (Künsch [11]), described in Appendix B. Figure 5a presents the results. The solid curve labelled "zero" gives the probability that the test statistic (i.e., $\hat{\zeta} - \hat{\xi}$) takes values less than zero. The dashed line gives the graph of the statistic itself to show its position relative to its own distribution. At higher bid levels, the probability that the difference is negative is very close to one, which represents a clear rejection of the theory.

Figure 6 plots estimates of $\widehat{\xi}$ and $\widehat{\zeta}$ for high l tracts. As in the case of low l tracts, $\widehat{\xi}$ is strictly increasing in bid and lies above the 45 degree line. Hence, Bayesian Nash equilibrium behavior is not rejected. The conditional expectation function $\widehat{\zeta}$ also lies above the 45 degree line, except at very low bid levels, and it is strictly increasing throughout. But, in contrast to the low l case, the difference $\widehat{\zeta} - \widehat{\xi}$ is close to zero at most bid levels. Figure 6a depicts a formal test of equality. The results indicate that the hypothesis of equilibrium bidding is not rejected at conventional confidence levels in the middle of the support of the bid distribution. There is some evidence of rejection at very low (i.e., less than \$500 thousand) and at high (i.e., more than \$10 million) bid levels.

An important feature of our tests is that we have not imposed any structure on the bidders' participation decisions in estimating ξ and ζ . We have simply assumed that expectations are consistent with the empirical law governing M_i . An alternative estimation strategy is to exploit the underlying structure of the probability law of M_i to estimate $G_{M|B}$ and ζ more efficiently using the entire vector of bids on each tract, including the "zeros", and estimates of the participation probabilities λ . We have not done so because we are not sufficiently confident of any specific participation model or our estimates of such a model to impose this structure. The main difficulty is with interpreting the "zeros". We would need to differentiate between potential bidders who did not bid because they were not active, and active bidders who chose not to bid because they did not obtain favorable information. We would also have to take a stand on whether bidders observe the number of active bidders. At this point, we prefer to forgo possible gains in efficiency in order to reduce the risk of misspecification.

6 Common vs. Private Value

In this section we discuss the problem of non-identification, or distinguishing between common and private valuation models, and whether the weight of the evidence favors one model over the other.

Laffont and Vuong [13] have observed that the CV model is not identified non-parametrically. Bids provide information about $w(s, s)$, but, given any increasing function $\psi(s)$,

$$w(s, s) = E[V|\psi(S) = \psi(s), \psi(Y) = \psi(s)].$$

Even if S is normalized so that $E[S|V = v] = v$, mean-preserving transformations of $F_{S|V}$ cannot be distinguished in the data. Hence, knowing the value of $w(s, s)$ at $b = \beta(s)$ is not sufficient to identify the value of s . Laffont and Vuong therefore argue

that any affiliated values model is observationally equivalent to some APV model if the reserve price is not binding.

The condition that the reserve price is not binding is important. In a private value environment, ξ must satisfy the boundary condition, $\lim_{b \downarrow r} \xi(b, G) = r$, which implies

$$\lim_{b \downarrow r} \frac{G_{M|B}(b|b)}{g_{M|B}(b|b)} \longrightarrow 0.$$

By contrast, in a common values environment, it follows from affiliation and the participation constraint that

$$\lim_{b \downarrow r} \xi(b, G) = E[V|S = s^*, Y = s^*] > r,$$

which implies

$$\lim_{b \downarrow r} \frac{G_{M|B}(b|b)}{g_{M|B}(b|b)} \longrightarrow c > 0.$$

Therefore, it is possible in principle to distinguish between the two models by examining the behavior of $G_{M|B}$ near the reserve price.⁸ Unfortunately, the reserve price in OCS auctions is not fixed. The government frequently rejected high bids near the announced minimum bid of \$15 per acre and, as a result, the bid distribution is too thin in this range to implement the test with any confidence. It might be possible to estimate the rejection decision of the government, i.e., the probability of accepting the high bid b given tract characteristics z , in conjunction with the empirical distribution of the highest rival bid, and then interpret G as the probability that a bid is high and is accepted.

An alternative approach to identification is to compute rents and bid markdowns under the alternative hypotheses of a private and common value model and determine which calculations yield more plausible numbers on economic grounds. Recall that, in the private values case, equation (3) can be interpreted as an inverse bid function. Thus, given our estimate of ξ , it is possible to infer bidders' private values from bid data under the assumption of Bayesian Nash equilibrium behavior. The fitted sample of valuations can then be used to calculate rents and markups. More precisely, let w_t denote the winning bid on tract t and define $\hat{\sigma}_{1:t} = \hat{\xi}(w_t)$ as the estimated private valuation of the winning bidder on tract t . Then the average value of rents under the

⁸A similar result holds for second-price auctions. Milgrom and Weber [18] observe that, in a CV environment, the difference between $E[V|S = s^*, Y \leq s^*]$ and the amount that s^* bids, $E[V|S_1 = s^*, Y_1 = s^*]$, implies a lower bound for the distribution of bids that is strictly above the reserve price, r . No such gap should be present in a private value environment.

private values hypothesis for a sample of tracts of size T is given by

$$R_\sigma = T^{-1} \sum_{t=1}^T [\hat{\sigma}_{1:t} - w_t].$$

Under the hypothesis of common values, the average value of rents is simply the difference between average net revenues and the winning bid. Let v_t denote the estimate of V on tract t . Then the average value of rents for a sample of tracts of size T is given by

$$R_v = T^{-1} \sum_{t=1}^T [v_t - w_t].$$

The values of R_σ and R_v for low l and high l tracts are reported in the table below. The standard deviations of the statistics are reported in parentheses. They are obtained from the bootstrap procedure that is described in the appendix.

	R_σ	R_v
Low l	\$19.0 (2.62)	3.75 (1.73)
High l	\$22.9 (2.50)	3.65 (2.89)

Private value rents vary between \$19 million and \$22.9 million on low l and high l tracts. Common value rents are much lower, approximately \$3.7 million on both sets of tracts. The latter fact suggests that bidders did not suffer from the “winner’s curse”, under the common values assumption.

Although we have referred to R_σ and R_v as rents, they are more accurately called quasi-rents since they do not include entry costs. To evaluate the plausibility of the figures, they should be compared to total entry costs on a tract. Under the assumption that entry rates are determined by a zero (expected) profit condition, the expected rent from bidding on a tract should be approximately equal to total entry costs. A bidder’s entry cost consists primarily of the price of a seismic survey, which is reputed to vary between \$9 and \$25 per acre, and the cost of hiring engineers to study the survey data and prepare a bid. The magnitude of these costs is probably several hundreds of thousands of dollars per bidder. Although the number of active bidders is not observed, even if we take the number of potential bidders as an upper bound, it is difficult to rationalize quasi-rents in the order of \$20 million. What type of costs prevent bidders from increasing their participation rates to take advantage of the available rents? By contrast, the quasi-rents calculated under the common value hypothesis seem more plausible. Given measurement error, one is unlikely to reject the hypothesis that the average rent from bidding is equal to total entry costs in either set of tracts.

The magnitudes of the markdown factors tell a similar story. Define $\hat{\sigma}_{it} = \hat{\xi}(b_{it})$ as the pseudo-value corresponding to bidder i 's bid on tract t . Under the hypothesis of private values, this is an estimate of bidder i 's valuation of tract t and the average value of the markdown can be defined as

$$M_{\sigma} = T^{-1} \sum_{t=1}^T \sum_{i=1}^{n_t} n_t^{-1} [\hat{\sigma}_{it} - b_{it}].$$

Under the assumption of common values, $\hat{\sigma}_{it}$ is an estimate of $E[V|b_{it}, M_{it} = b_{it}]$. Consequently, M_{σ} can also be interpreted as an estimate of a markdown but, unlike the private values case, it is not a measure of the amount that a bidder expects to earn in the event of winning. The latter is determined by conditional expectation $E[V|b, M < b]$. Define \hat{w}_{it} as the estimate of bidder i 's valuation of the tract conditional upon winning with a bid of b_{it} . It is obtained by evaluating the function $E[V|w = b]$, shown in Figures 1 and 2 for low l and high l tracts respectively, at $b = b_{it}$. The average markdown percentage under common values is then

$$M_v = T^{-1} \sum_{t=1}^T \sum_{i=1}^{n_t} n_t^{-1} [\hat{w}_{it} - b_{it}].$$

The values of M_{σ} and M_v and their standard deviations are reported below.

	M_{σ}	M_v
Low l	14.65 (1.67)	3.76 (1.54)
High l	14.43 (1.30)	3.56 (1.30)

Under the private values hypothesis, bidders marked down their bids independently of the number of potential bidders and by slightly more than \$14 million. Under the alternative hypothesis of common values, the markdown is roughly the same on both sets of tract and equal to rents. Thus, there is no adverse selection associated with winning, which is consistent with bidders anticipating the “winner’s curse”. We also calculated the average markdown as a percentage of average value. Under private values, the percentage markdowns at the average are 88% for low l tracts and 72% for high l tracts. In other words, firms are bidding slightly more than 1/9 of their estimated of the bidder specific value on low l tracts, and less than 1/3 on high l tracts. These figures are lower than the percentage markdowns calculated by Li, Perrigne and Vuong [14], who consider only the subset of tracts with two or three bids and obtain an estimate of 1/3. The corresponding percentage markdowns under common values are 65% and 38% respectively. The latter numbers seem more plausible.

An additional difficulty with assuming private values alone is that there is considerable within tract variation in bids (Porter [20]). For example, the second highest bid averages about 56 per cent of the highest bid. In a private values model, such extreme variation in bids can be rationalized only if there is large variation in the private components of values. Yet the variation in valuations necessary to generate the observed variation in bids is an implausible order of magnitude.

In summary, the OCS lease auctions probably contain both common and private valuation components, and hence are best described by a general affiliated values model, but the data suggest that the former components are more important.

7 Conclusion

The main conclusions to be drawn from the preceding empirical analysis are (1) that the “winner’s curse” is evident in the data and (2) that the bidders are aware of its presence and bid accordingly. Bidding behavior appears to be largely consistent with a symmetric common value environment, given our measure of potential competition and our measure of ex post returns. In contrast, some features of bidding behavior appear to be inconsistent with a private values environment. To repeat, the bidders’ valuations in OCS auctions probably have both private and common components, but the common components appear to be important.

The theoretical model of the paper may be employed to examine the participation decisions of the potential bidders, and thereby obtain more efficient estimates of bidding rules. We hope to pursue this direction in future work.

References

- [1] Bajari, P. (1998): “Econometrics of the First Price Auction with Asymmetric Bidders,” mimeo, Harvard University.
- [2] Bajari, P. (1998): “Econometrics of Sealed Bid Auctions,” *Proceedings of the Business and Economic Statistics Section of the American Statistical Association*, 41-49.
- [3] Baldwin, L., R. Marshall and J.F. Richard (1997): “Bidder Collusion at Forest Service Timber Sales,” *Journal of Political Economy*, 105, 657-699.
- [4] Donald, S., and H. Paarsch (1993): “Piecewise Maximum Likelihood Estimation in Empirical Models of Auctions,” *International Economic Review*, 34, 121-148.
- [5] Elyakime, B., J.J. Laffont, P. Loisel and Q. Vuong (1994): “First-Price, Sealed-Bid Auctions with Secret Reservation Prices,” *Annales d’Economie et de Statistique*, 34, 115-141.
- [6] Hall, P. (1992): *The Bootstrap and Edgeworth Expansion*, New York: Springer.
- [7] Hendricks, K., and R. Porter (1988): “An Empirical Study of An Auction with Asymmetric Information,” *American Economic Review*, 78, 865-883.
- [8] Hendricks, K., R. Porter and C. Wilson (1994): “Auctions For Oil and Gas Leases with an Informed Bidder and a Random Reservation Price,” *Econometrica*, 62, 1415-1444.
- [9] Hong, H., and M. Shum (1999): ”Structural Estimation of First Price Auction Models: Measuring Common Values and the Winner’s Curse,” mimeo, Princeton University.
- [10] Klemperer, P. (1998): ”Auctions with Almost Common Values: The ‘Wallet game’ and its Applications,” *European Economic Review*, 42, 757-769.
- [11] Künsch, H.R. (1989): “The Jackknife and the Bootstrap for General Stationary Observations,” *Annals of Statistics*, 17, 1217–1241.
- [12] Laffont, J.J., H. Ossard and Q. Vuong (1995): “Econometrics of First-Price Auctions,” *Econometrica*, 63, 953-980.
- [13] Laffont, J.J., and Q. Vuong (1996): “Structural Analysis of Auction Data,” *American Economic Review, Papers and Proceedings*, 86, 414-420.

- [14] Li, Tong, I. Perrigne and Q. Vuong (1998): “Affiliated Private Values in OCS Wildcat Auctions”, mimeo, University of Southern California.
- [15] Marshall, R., and M. Raiff (1997): “The Impact of Synergies on Bidding in the Georgia School Milk Market,” mimeo, Pennsylvania State University.
- [16] McAfee, R.P., D. Quan, and D. R. Vincent (1996): “How to Set Minimum Acceptable Bids, with an Application to Real Estate Auctions,” mimeo.
- [17] McAfee, P., and D. Vincent (1992): “Updating the Reserve Price in Common-Value Auctions,” *American Economic Review, Papers and Proceedings*, 82, 512-518.
- [18] Milgrom, P., and R. Weber (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089-1122.
- [19] Paarsch, H. (1992): “Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions,” *Journal of Econometrics*, 51, 191-215.
- [20] Porter, R. (1995): “The Role of Information in U.S. Offshore Oil and Gas Lease Auctions,” *Econometrica*, 63, 1-27.
- [21] Smiley, A. (1979): *Competitive Bidding Under Uncertainty: The Case of Offshore Oil*, Ballinger, Cambridge.
- [22] Stone, C.J. (1977): “Consistent Nonparametric Regression,” *Annals of Statistics* 5, 595–645.
- [23] Wilson, R. (1977): “A Bidding Model of Perfect Competition,” *Review of Economic Studies*, 44, 511-518.

Appendix

A Estimation Methods

There are four basic functions that need to be estimated: $E[V|B = b]$, $E[V|B = b, M < b]$, $E[V|B = b, M = b]$, and $\xi(b, G)$. We estimate each of these functions nonparametrically. The procedure is described below. All computations employ log (base 10) bids.

A.1 Univariate Local Linear Regression

We estimate the conditional expectation function $E[V|B = b]$ on the sample of all bids $\{b_{it}\}$ (within the same number of potential bidders class) and tract values $\{v_t\}$ using the following formula at each value of b (Stone [22])

$$\{\hat{v}(b), \hat{r}_v(b)\} = \arg \min_{v,r} \sum_{t=1}^T \sum_{i=1}^{n_t} \{v_t - v - r(b - b_{it})\}^2 k\left(\frac{b - b_{it}}{h}\right) \frac{1}{n_t}, \quad (6)$$

where T is the number of tracts in the sample, h is the *bandwidth* and k the *kernel*. We use a standard normal kernel and a bandwidth chosen by visual inspection.

The conditional expectation $E[V|B = b, M < b]$ is estimated on the sample of winning bids $\{w_t\}$ and tract values $\{v_t\}$ using the formula,

$$\{\hat{w}(b), \hat{r}_w(b)\} = \arg \min_{v,r} \sum_{t=1}^T \{v_t - v - r(b - w_t)\}^2 k\left(\frac{b - w_{it}}{h}\right). \quad (7)$$

Each term in the summation of equation (6) includes a weight equal to the reciprocal of the number of bids on tract. This weighting is necessary to compare $E[V|B = b, M < b]$ and the other statistics, since only one bid per tract is used for the computation of $E[V|B = b, M < b]$, whereas all bids are used in the computation of the other statistics.

A.2 Bivariate Local Linear Regression

The conditional expectation $E[V|B = b, M = b]$ is estimated on the sample of all bid pairs $\{b_{it}, m_{it}\}$ and tract values $\{v_t\}$ (within the same number of potential bidders

class) using the formula

$$\begin{aligned} & \{\widehat{v}(b, m), \widehat{r}(b, m)\} \\ &= \arg \min_{v, r_1, r_2} \sum_{t=1}^T \sum_{i=1}^{n_t} \\ & \quad \{v_t - v - r_1(b - b_{it}) - r_2(m - m_{it})\}^2 k\left(\frac{b - b_{it}}{h_b}\right) k\left(\frac{m - m_{it}}{h_m}\right) \frac{1}{n_t}, \end{aligned} \quad (8)$$

where h_b and h_m are bandwidths.

A.3 Pseudo-Values

As in Li, Perrigne and Vuong [14], we estimate G_L and g_L , the distribution and density of base 10 log bids, and then derive the ratio of the distribution and the density of bid levels according to $G(b|b)/g(b|b) = b \log(G_L/g_L)$. Note that $G(b|b) = P[M < b | B = b] = E[I(M < b) | B = b]$, where I is an index function. We estimate $G(b|b)/g(b|b)$ using the formula

$$\frac{h_m \sum_{t=1}^T \sum_{i=1}^{n_t} k([b - b_{it}]/h_b) I(m_{it} < b) n_t^{-1}}{\sum_{t=1}^T \sum_{i=1}^{n_t} k\left(\frac{b - b_{it}}{h_b}\right) k\left(\frac{m - m_{it}}{h_m}\right) \frac{1}{n_t}}. \quad (9)$$

B Bootstrap

The bootstrap procedure used is a spatial block bootstrap. Except for the rent calculations, the number of bootstrap replications is 10,000. For the rent calculations, the estimated standard deviations are based on 1,000 bootstrap replications. The computations were carried out on a 300 MHz PowerMac running Mac OS X Server.

The bootstrap procedure is not pivotal and the bootstrap approximation error therefore decreases at the same rate as when the asymptotic distribution is derived by first order asymptotics (Hall [6]).

The procedure is outlined below.

1. Number all tracts 1 through T .
2. Denote by q_t the reciprocal of the number of tracts in Λ_t , the neighborhood of tract t , including tract t itself.
3. Let $Q_t = \sum_{s=1}^t q_s / \sum_{s=1}^T q_s$
4. Draw a random number r in $[0, 1]$.

5. Find the tract t whose value of Q_t is the closest to r .
6. Add tract t to the bootstrap sample.
7. Add all tracts in the neighborhood of tract t to the bootstrap sample.
8. Repeat steps 4–7 until the number of tracts in the bootstrap sample is greater than or equal to T .
9. Eliminate tracts at the end of the bootstrap sample such as to make the number of bootstrap tracts equal to the number of sample tracts.
10. Compute all statistics for the bootstrap sample.
11. Repeat steps 4 through 10 R times, where R denotes the number of replications.

Several points about the procedure are worth noting. First, the block bootstrap is needed to accommodate dependence across tracts. In the ‘normal’ bootstrap one would make T independent draws from the tract distribution. We make a number of independent draws from the neighborhood distribution, such as to make the number of tracts in the bootstrap sample equal to T . Second, because a tract is selected either directly or as part of the neighborhood of a directly selected tract, tracts in regions where a lot of tracts are bid are more likely to be selected. This is countered in step 5 by making the probability that tract t is selected directly inversely proportional to the number of tracts in Λ_t . The procedure is not perfect, but it eliminates most of the variability. Finally, the number of bids varies a little from bootstrap sample to bootstrap sample, but generally by less than 5% in either direction.

TABLE 1**Wildcat Bidding by the Twelve Firms and Consortia with the Most Bids**

Firms and Consortia	Number of Solo Bids	Number of Joint Bids	Potential Bidder	Participation Rate
Arco/Getty/Cities/Cont.	437	114	1027	0.54
Standard Oil of California	408	76	1022	0.47
Standard Oil of Indiana	132	276	905	0.45
Shell Oil	444	3	981	0.46
Gulf Oil	201	81	801	0.35
Exxon	325	42	812	0.45
Texaco	114	178	823	0.35
Mobil	48	163	700	0.30
Union Oil of California	95	201	805	0.37
Phillips	98	65	498	0.33
Sun Oil	241	93	723	0.46
Forest	195	0	493	0.40

TABLE 2**Summary of Wildcat Sales, 1954-1970**

Sale Date	#Tracts Offered	#Tracts Bid	#Tracts Big 12 Bid	#Tracts Sold	#Tracts Drilled	#Hits	Mean Rev	Mean NetRev	Mean Hibid
54-10-13	199	90	77	90	65	45	44.00	6.99	4.49
54-11-09	38	19	17	19	10	4	13.38	0.46	4.27
55-07-12	210	117	92	117	64	27	30.62	2.23	3.14
60-02-24	385	173	141	147	117	61	89.25	21.61	4.97
62-03-13	401	211	169	206	165	79	56.57	11.17	2.47
62-03-16	410	210	169	205	169	79	52.59	9.69	3.75
67-06-13	206	172	142	158	130	53	67.09	10.55	7.80
68-05-21	169	141	110	110	71	16	12.21	-0.87	10.72
70-12-15	127	127	57	119	112	64	68.76	19.60	15.18
Total	2145	1260	974	1171	903	428	58.60	10.48	6.19

*Dollar figures in millions of 1982 dollars.

TABLE 3**Characteristics of Wildcat Tracts and Their Neighborhoods by Number of Potential Bidders*****Number of Potential Bidders (I)**

	0	1	2	3	4	5	6	7	8	9	10	11	12	Total
# of Tracts	7	45	74	115	72	94	101	138	164	189	160	85	16	1260
Bids/Tract (All)	1.00	1.13	1.55	1.99	2.17	2.84	2.80	3.94	3.65	5.12	5.11	5.36	4.25	3.62
Bids/Tract (12)	0.00	0.91	1.24	1.53	1.64	2.33	2.32	3.01	2.65	3.77	3.51	3.66	2.94	2.67
Tract Hibid	0.60	1.32	2.21	2.63	2.78	4.01	2.94	6.79	5.70	9.26	11.5	8.59	13.1	6.19
	(0.38)	(1.72)	(3.82)	(4.38)	(4.08)	(6.63)	(3.62)	(11.3)	(9.36)	(13.5)	(18.8)	(12.5)	(22.6)	(11.6)
% Tract Drilled	66.7	62.2	54.9	71.8	73.9	79.3	73.2	81.4	79.6	82.3	80.0	84.2	100.0	77.1
% Hits	25.0	39.1	38.5	45.6	41.7	50.8	50.7	47.6	41.3	47.7	54.2	54.7	46.7	47.4
Tract	3.60	84.8	28.0	66.1	60.6	38.2	45.2	76.8	43.8	61.1	71.5	55.8	62.7	58.6
Revenue	(0.0)	(80.3)	(41.4)	(101)	(89.4)	(67.4)	(81.7)	(102)	(73.4)	(97.5)	(103)	(66.6)	(89.2)	(89.8)
Tract	-1.30	10.91	2.45	11.3	10.9	5.97	7.53	15.4	6.02	11.7	15.5	12.5	15.4	10.5
NetRev	(1.31)	(33.4)	(15.7)	(41.8)	(39.1)	(29.5)	(36.2)	(49.3)	(31.3)	(49.0)	(49.8)	(34.8)	(42.2)	(40.9)

*Dollar figures in millions of 1982 dollars, and refers to the means of the preceding sample. Standard deviations are in parentheses.

Table 4: Bid Frequency Distribution

