A Bidding Model for Auctions of Offshore Alternative Energy Sites

prepared for

The Eastern Economic Association
New York City
March 1, 2009

by

Radford Schantz
US Department of the Interior
381 Elden Street
Herndon, VA
Radford.schantz@mms.gov

Walter Stromquist
Department of Mathematics
Swarthmore College
Swarthmore, PA 19081
mail@walterstromquist.com

Abstract. This paper presents an auction model that predicts the probability of at least one bid, the number of bids, and the high bid for rights to develop a resource such as wind energy at a given location. Factors that tend to increase the potential number of bidders include the bidders' valuations of the site, the availability and cost of information, and minimum bid.

Currently applied to offshore oil lease auctions, where there is uncertainty about the amount of oil underground, it is proposed that the model be extended to cover alternative energy in the offshore context, where there is uncertainty about the wind energy at particular sites. Generally, uncertainty means that an auction tends to award the rights on the basis of an over-estimate of site value. This "winner's curse" is less significant in an ascending auction where bidders can learn by observing other bids.

In our model, resource uncertainty can be reduced by acquiring more information. Information is of two kinds, public and private, and there are two types of bidders, those with both types of information and those on the fringe with only public information. Each potential bidder buys information with a certain probability, and the equilibrium conditions determine the number of bidders with private information. There can be bidders also who chose to rely on only free, public information. Bidding strategy is derived from a starting condition and a differential equation. From the bidding strategy, the probability of a bid, the number of bids (probability of competition), and the winning bid amount can be inferred.

Introduction

We describe a model of first-price, sealed-bid auctions, based on a common value assumption. The model is being used by the US Department of the Interior for analyzing auctions of offshore alternative-energy sites.

The Interior Department, and more particularly the bureau called Minerals Management Service, have for decades managed the Federal program for leasing offshore oil and gas resources. This activity has included regular lease auctions that have, over the history of the activity, resulted in well over $60 billion in bonus revenue for the government and a larger total of royalty and tax revenue. Onshore leasing is managed by a different bureau, the Bureau of Land Management.

There is now interest in leasing sites for alternative energy production --- including sites suitable for wind, solar, geothermal, and ocean-current ("hydrokinetic") energy. As in the case of oil and gas leases, the offshore part of the activity is managed by the MMS.

Auctions of alternative-energy sites raise some issues that are different from the oil-and-gas program. Even to the extent that the issues are similar, the new focus offers a chance to re-evaluate various auction formats. A versatile bidding model will be helpful in this analysis. The purpose of this paper is to describe such a model. We hope to gain from public comment.

Our model is an extension of one that is currently implemented as part of IMODEL, a computer program currently used by the MMS for forecasting offshore oil-and-gas lease auctions and for policy analysis [1]. In extending the model, we have benefited from reading a paper by Thompson and Wright [2]. While the bidding model in [2] is designed for a more limited research purpose, it includes an information model (private signals with independent additive errors) on which we have drawn heavily.

The first sections of this paper describe the offshore program for alternative-energy leasing. The model itself is introduced in subsequent sections. A simple example is presented along with the model itself, mainly to clarify the concepts involved. A more detailed description appears in the appendix.

Oil and gas auctions

Two different auction systems in use at the Interior Department for oil and gas leases: first-price sealed-bid, for offshore, and open ascending second-bid, for onshore.

Auctions are required by law when leasing offshore oil and gas resources. The main authority is the Outer Continental Shelf Lands Act (1953), 43 USC 1331, which provides no process for noncompetitive award. In the case of onshore resources, auctions are also used when there is competitive interest. Absent that interest, an onshore lease can be awarded noncompetitively. The main authority is the Federal Oil and Gas Leasing Reform Act (1987).

Both onshore and offshore leases can be resold and traded after the auction; so the secondary market serves to adjust the allocation of leases to companies. The auctions are repeated, e.g. annually for most of the Gulf of Mexico.

Alternative energy site auctions

Alternative energy for this paper includes wind, solar, wave, ocean current, and geothermal. As there are numerous types of alternative energy and correspondingly numerous licensing regimes, let us focus temporarily on wind energy sites.

Onshore, BLM administers wind energy rights-of-way (ROW). The main authority is the Federal Land Policy and Management Act (1976), 43 USC 1701. There are three basic types of ROW: energy testing and monitoring; site testing and monitoring; and commercial development. For onshore sites, applications are sometimes handled on a first-come, first-serve noncompetitive process. However, BLM will use a competitive process if there is a land use plan that requires it or if more than one applicant shows substantive interest.

One interesting instance of a BLM wind site auction is the Palm Springs Section 28 auction held in 2004. The site covered 285 acres that had been occupied formerly by an older wind farm. BLM advertised for expression of interest in repowering this abandoned wind energy site. Four qualified companies responded.

An interesting feature of this auction is the subtle definition of the item being offered. The purpose of the auction was to award a preference right to apply for an ROW in the future; the present auction was, that is to say, to select an applicant to develop the site eventually (if environmental hurdles, etc., were overcome. After considering, in the way of possible competitive processes, and rejecting lottery and multi-factor subjective assessment, the open ascending auction format was adopted.

To begin, persons had 45 days to submit written bids with deposits. After that, bidders could submit higher bids in writing or by fax. The latest high bid was posted on the internet, keeping bidder identifies confidential. The increment was $500. If bids were still being submitted on the 45th day, bidding would continue for another day, and it would continue day by day until higher bids stopped. In fact, the auction that began in June ended in November – five months. The winning bid was $250,000.

Offshore alternative energy

Section 388 of the Energy Policy Act of 2005 amended the OCS Lands Act and gave Interior authority to grant leases for alternative energy activities. Competition is provided for in the following manner:

> COMPETITIVE OR NONCOMPETITIVE BASIS- …the Secretary shall issue a lease, easement, or right-of- way under paragraph (1) on a competitive basis unless the Secretary determines after public notice of a proposed lease, easement, or right-of-way that there is no competitive interest.

That public notice is like a tract nomination step, and it helps us identify the area to be offered. In general, we anticipate that there will actually be instances of both kinds of award. As for onshore, the offshore auctions might take various forms depending on the type of lease, and we are open to advice.

To begin implementation of this Act, a proposed rule, titled Alternative Energy and Alternate Uses of Existing Facilities on the Outer Continental Shelf, was published in 2008. Final publication of the rule is still pending. Meanwhile, under an interim policy, Interior called for information and area nominations for offshore measurement leases. These will be for five years or more, but they cannot be converted into commercial leases. Forty nominations were received, and we are proceeding with leasing in 16 areas that we consider top priority. They are:

New Jersey	6 wind	Noncompetitive
Delaware	1 wind	Noncompetitive
Georgia	3 wind	Noncompetitive
Florida	4 current	3 competitive, 1 noncompetitive
California	2 wave	1 competitive, 1 noncompetitive

Figure 1 is a map showing the New Jersey areas. Figure 2 illustrates three actual or proposed alternative energy projects for offshore.

Figure 1. New Jersey wind energy sites

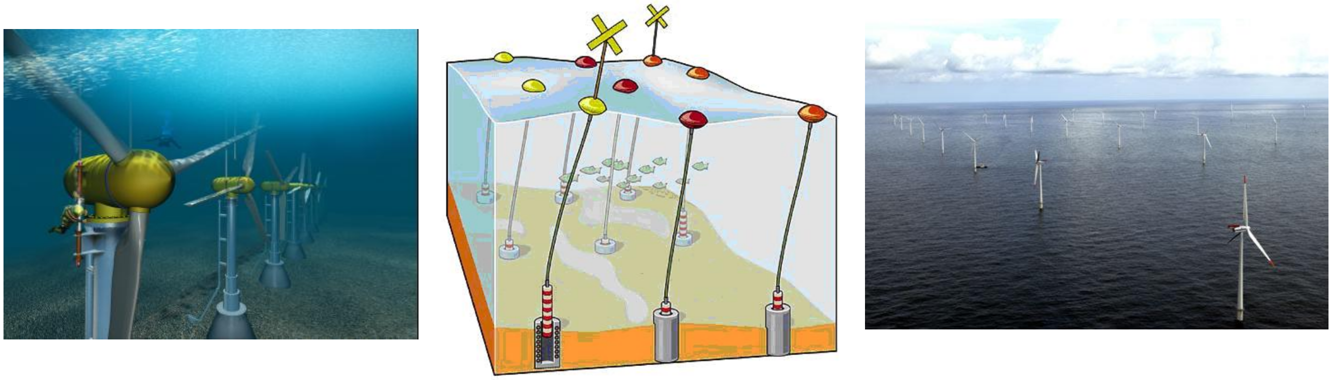

Figure 2. Alternative energy possibilities

The Auction Model -- Overview

The auction model describes a first-price, sealed-bid auction. It incorporates a common value assumption: that is, the value V of the property being offered is the same for all participants, and bids differ only because bidders have different estimates of V. The model treats the auction as a non-cooperative game. It calculates a symmetric equilibrium bidding strategy for the participants. The principal outputs from the model are the expected number of bids and the expected revenue for the seller.

The model has several features worthy of special attention. The number N of potential bidders is itself a random variable; the user may either specify a mean value for N, or allow the distribution to be determined internally based on an assumed information cost and an equilibrium condition. The possible role of uninformed bidders is modeled explicitly. Private signals are modeled in a flexible and realistic way.

The model is specified in full detail in the appendix. The next few sections contain a less detailed overview, along with a simple example intended mainly to illustrate the concepts.
The example is typical of oil-and-gas properties in the Gulf of Mexico (to the extent that a single example can be typical).

The Auction Model, in more detail

a. Rules of the auction. The seller announces a minimum bid, b_{min}, and (optionally) a bidding fee. Anyone may submit a bid, provided that it is at least equal to b_{min}. Each bidder must pay the bidding fee, regardless of the outcome. Bidders submit sealed bids, without knowing how many other bids are being submitted. The highest bidder receives the property and pays the amount of its own bid.

In our example, we take the minimum bid to be $128 thousand (typical of oil-lease auctions) and the fee to be zero. (There is no provision in current law for a fee; it is an extra capability of the model only because it is not to difficult to include.)

b. The common value. The property has a value, V, which is the same for all bidders. None of the participants knows the exact value of V, although they have both public and private information on which to base estimates.

We model V as a random variable. More precisely, we define another random variable, U, which we call the *value parameter*, and which is related to V by the (increasing) function

$$V = v(U).$$

The reason for this extra level of abstraction is that we want the private signals to be subject to independent, additive errors. But people generally make larger errors (in dollars) when estimating larger amounts. That means that if the bidders are estimating V directly, it is not plausible that their estimating errors are independent of V. It is more plausible that the estimation errors are independent of some transformed version of V, and that is the role played by U. We therefore assume that the bidders are estimating U, and we allow an arbitrary (monotone) relationship between V and U.

We assume that all available public information about V (or U) is represented by a *prior distribution* on U. (This implies a prior distribution on V, but we do not need to represent that distribution explicitly.) The model allows an arbitrary distribution H. That is, H is a cumulative distribution function (cdf) for U, defined by

$$H(u) = \Pr\left(U \le u\right).$$

This distribution may be continuous, discrete, or some mixture of these forms.

The prior distribution H and the value function v are assumed to be common knowledge.

In our example, we take H to be a discrete distribution, as follows:

With probability P_1 = 0.20, we have U = +1/2 and V = $1 million.
(This represents the possibility that the property is economically successful.)

With probability P_0 = 0.80, we have U = -1/2 and V = - $ 90 thousand.
(This represents the possibility that the property is a failure, and that the operator must invest additional money before finding that out.)

The graph of H is shown below. We would normally present a graph of the function v, but since there are only two possible values for U, only two values of v(u) need to be specified.

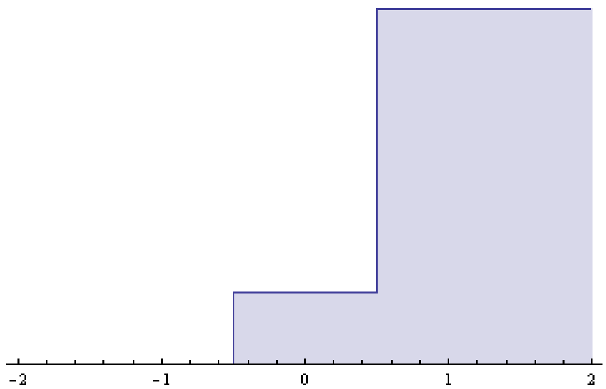

It follows from these assumptions that the prior mean value of the property (based on public information) is

(0.20) ($ 1 million) + (0.80)(–$ 90 thousand) = $ 128 thousand,

which is barely enough to justify a minimum bid. In this example we should expect that a large fraction of potential bidders will choose not to bid.

c. How many bidders? We now introduce a bid of jargon. The potential bidders are called *evaluators.* We reserve the term "bidders" for evaluators who actually choose to bid.

We suppose that there are a very large number of evaluators, and that a few of them are selected at random to receive *private signals* and so to become *informed evaluators.* The number of informed evaluators is a random variable, N.

In the simplest version of the model, N is a Poisson random variable, with mean m specified by the user. This means that

$$\Pr(N = n) = q(n) = \frac{m^n}{n!} e^{-m}$$

for n = 0, 1, 2, …. In particular, there is a positive probability that N = 0.

The model permits an alternative. The user may specify an information cost, k, in dollars. Evaluators then decide independently whether to pay k, and to become informed evaluators. At equilibrium, this results in a Poisson distribution for N as above, but with m determined inside the model as a function of k. (This makes no difference for modeling single properties. Since specifying k determines a unique value of m, the user might as well specify m. When modeling a range of cases with different distributions for V, however, it may make more sense to keep k fixed across cases than to keep m fixed. Keeping k fixed means that the tracts that are more valuable, *a priori,* attract more informed evaluators. This matches experience.)

In our example, we fix m = 3.0.

This is the mean number of informed evaluators. The actual number of bidders depends on the equilibrium strategy; as some of the evaluators --- the ones with discouraging private signals --- will choose not to bid. In fact, given the borderline *a priori* value, even evaluators with fairly neutral private signals will choose not to bid. In this example, we should expect that only evaluators with encouraging signals will be able to justify a minimum bid.

d. What about uninformed evaluators? Anyone is allowed to bid, even evaluators who do not receive private signals. In most cases it is not rational for uninformed evaluators to bid, so they do not influence the outcome of the auction. That is the case with our example.

But there are cases in which uninformed evaluators play a role. For example, suppose that the lowest possible value of V (according to the public prior distribution) is $200 thousand but the minimum bid is $128 thousand. Then it is rational for uninformed evaluators to bid up to $200 thousand --- they can't lose --- and it is reasonable to assume that somebody, at least, will bid that amount. That means that the strategy for informed evaluators has to start at $200 thousand, not at $128 thousand. In effect, the uninformed evaluators are enforcing a more realistic starting point than the announced minimum bid. (The starting point becomes an initial condition for a differential equation.)

The model solves for the optimal behavior of uninformed evaluators as part of its equilibrium calculation.

Since, in our example, uninformed evaluators do not bid, we will have no more to say about uninformed evaluators outside of the appendix.

e. Private signals. Each informed evaluator receives a *private signal*, which is a number that depends probabilistically on the value parameter U. If X_i is the i-th evaluator's private signal, then X_i is modeled as a random variable. It is given by

$$X_i = U + R_i$$

where R_i is also a random variable, the i-th evaluator's *estimation error*. We assume that the estimation errors are independent and identically distributed, and that they are independent of both U and N.

We assume, further, that the estimation errors are normally distributed, with a common standard deviation σ for all evaluators. The standard deviation is a model parameter, and we prefer to specify it in terms of the *information parameter* $c = 1/\sigma^2$ (the inverse variance). It is fair to think of c as an "amount of information" available to each informed evaluator.

This is the same model used by Thompson and Wright. (They allow non-normal distributions for R_i. For their purpose it suffices to regard X_i as an estimate of V directly. We treat X_i as an estimate of U.)

The user specifies c, which determines σ. The error model and the values of c and σ are assumed to be common knowledge. Each i-th evaluator knows the value of X_i, but not R_i or the value of any other evaluator's private signal. We can now write the density function of each estimation error:

$$f(r) = \sqrt{\frac{c}{2\pi}}\, e^{-cr^2/2}$$

We use F(r) to denote the corresponding cumulative distribution function (cdf).

In the example we take c = 0.25 (so that σ = 2.0). This is realistic for oil and gas tracts, when we use the particular prior that we have chosen for the example.

The two-point prior makes it easier for us to illustrate the information model. Since the only uncertainty in V is whether the tract is a success or failure, each evaluator's private signal is used as a way to try to distinguish the two cases.

Figure 3 illustrates the distribution of private signals. Although U can only take on two possible values (-1/2 and +1/2), a private signal X can take on any real value. Most signals (regardless of U) are between -6 and +6. The right-hand curve is the distribution of private signals for tracts that are actually successes, and the left-hand curve is the distribution for tracts that are actually failures.

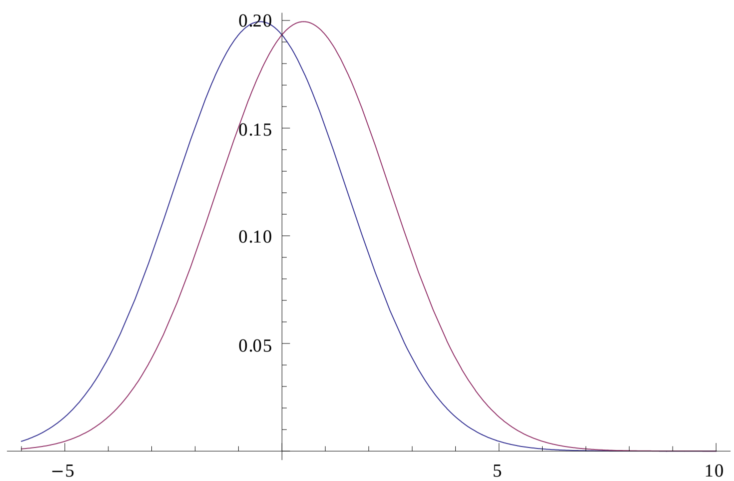

Figure 3. Density of private signals
(Left curve if tract is actually a failure, right curve if it is actually a success)

Clearly a private signal gives the evaluator some help in identifying the successful tracts, but not with very much help.

We can compute the posterior probability of success, given a private signal x. The prior probability is $P_1 = 0.20$, so the posterior probability is

$$\frac{(0.20)\,f(x-1/2)}{(0.20)\,f(x-1/2)+(0.80)\,f(x+1/2)}$$

which turns out to be

Posterior probability of success = $1/(1+4e^{-cx})$.

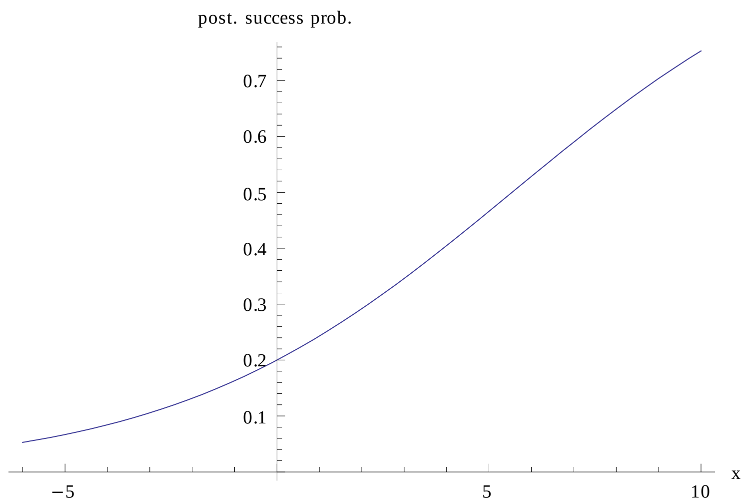

Figure 4. Posterior probability of success, as a function of the private signal x

This is a logistic curve. Note that it takes a private signal of about +6 (rare and extremely positive) to raise the posterior probability even to 1/2. Most private signals just permit the evaluator to make a slight adjustment to the available public information. (This is reality.)

On the other hand, a small edge in estimating the success probability can have a large effect on the estimated value of the property. A private signal of $x = 1.67814$, for example, increases the posterior probability of success to just 0.2755, but that corresponds to a posterior value estimate of $210 thousand. As we will see, a private signal of $x = 1.67814$ justifies a minimum bid even in the face of the winners' curse.

Of course, this calculation is not directly relevant to an evaluator's bidding strategy. The evaluator can also draw inferences from the mere fact of being an evaluator. (For example, N can't be zero!) Further, as we know from the theory of auctions, the evaluator should be estimating the value conditioned on his signal being x, but also on that being the highest signal. (That's how to avoid the winner's curse.) For some purposes, other conditioning is required. The calculation we have just seen is just an illustration of the information model.

f. Summary of the inputs. The model has these input parameters:

b_{min}, fee (rules of the auction)

H	(cdf for the value parameter)
v	(function defining value, V = v(U))
m	(the mean number of informed evaluators)
c	(information parameter; determines sigma, f, F)

All of these are assumed to be common knowledge.

g. What about the Winner's Curse?

We model this auction as a non-cooperative game. We will find an equilibrium strategy, which consists mainly of a function g(x) that tells an evaluator how much to bid, if its private signal is x.

If g(x) were equal to the estimator's posterior estimate of V, we would expect the evaluators to take huge losses on average, because of the winner's curse. That is, the evaluator would tend to win only those auctions in which it had seriously overestimated the value, and it would pay a price for those overestimates, with no compensation for having underestimated in other cases.

Instead, we should expect g(x) to be much less than the posterior estimate of value (for each value of x), for two reasons. This is how bidders compensate for the winner's curse, and this is how they seek to make a profit from the happy circumstance of being informed evaluators.

h. The symmetric equilibrium

These assumptions are enough to determine the equilibrium strategies of all participants.

A symmetric equilibrium consists of a number x* and a function g( ). The number x* is the smallest private signal that justifies a minimum bid. The function g(x) gives the optimum bid for an evaluator whose private signal is x. Note that

$$g(x^*) = b_{min},$$

and that g(x) is defined only for $x \geq x^*$. When it is defined, we assume that g( ) is strictly increasing and has a positive derivative.

(We are continuing to disregard the role of uninformed evaluators. In this model we never consider asymmetric equilibria.)

Calculation of the equilibrium strategy is far from trivial, but it can be done using standard methods. The function g arises finally as the unique solution to a differential equation with initial condition $g(x^*) = b_{min}$. The calculation is done numerically in the model, and x* is found be a search. (Analytic solutions are possible for some choices of H and v, but if we want full generality in the inputs, there is no alternative to numerical solution.)

Let's compute the equilibrium strategy in our example. First, here is a graph of the expected profit of an evaluator who bids the minimum bid, as a function of the private signal x. (That is, this shows the evaluator's posterior estimate of its profit, conditional on all of the circumstances, including the signal x, and the equilibrium behavior of other participants.)

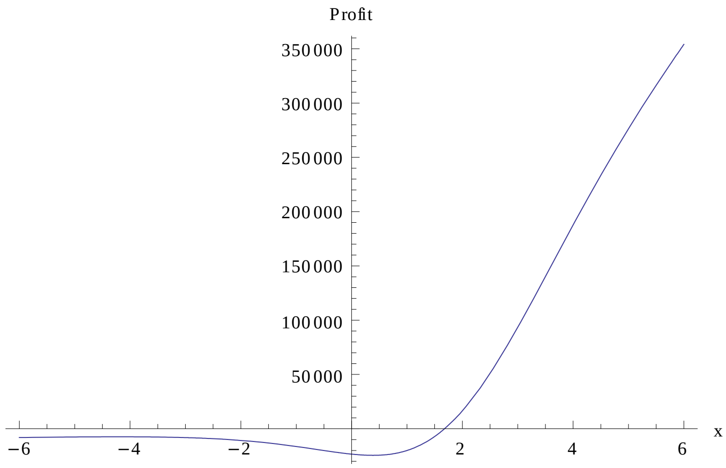

Figure 5. Expected profit resulting from bidding the minimum bid,
from the point of view of an evaluator with private signal x.
(x is on the horizontal axis, and the expected profit is in dollars.)

There is a minimum value x* that justifies the minimum bid, and it is the point at which this curve crosses the x axis. It is found by solving the underlying equation numerically. In this example, it is

$$x^* = 1.67814.$$

That is an encouraging signal, but not unusually so.

If the property is actually a success, the probability of any particular evaluator getting this signal or better is

$$1 – F (x^* - 1/2) = 0.278.$$

Since the expected number of evaluators is m = 3, it is far from certain that any evaluator will bid, even if the tract is actually a success. If the tract is a failure, the corresponding probability is

$$1 – F (x^* + 1/2) = 0.138.$$

We can already calculate the unconditional probability that at least one evaluator will bid. With probability 0.20 it is

$$1 - \exp(- 3 (0.278)) = 0.566$$

and with probability 0.80 it is

$$1 – \exp(-3 (0.138)) = 0.339,$$

so with no condition it is

$$(0.20) (0.566) + (0.80) (0.339) = 0.384.$$

The model has permitted us to conclude (a) that the probability of getting at least one bid on this tract is 0.384, and (b) that tracts that are actually successes are more likely to draw bids than those that are actually failures. This is encouraging.

Now we compute the equilibrium strategy in the example. (The heavy calculations presented here are done using Mathematica. This particular calculation results in an “interpolating function,” which means that we have resorted to numerical methods.) Here is a graph of the equilibrium bidding function g(x):

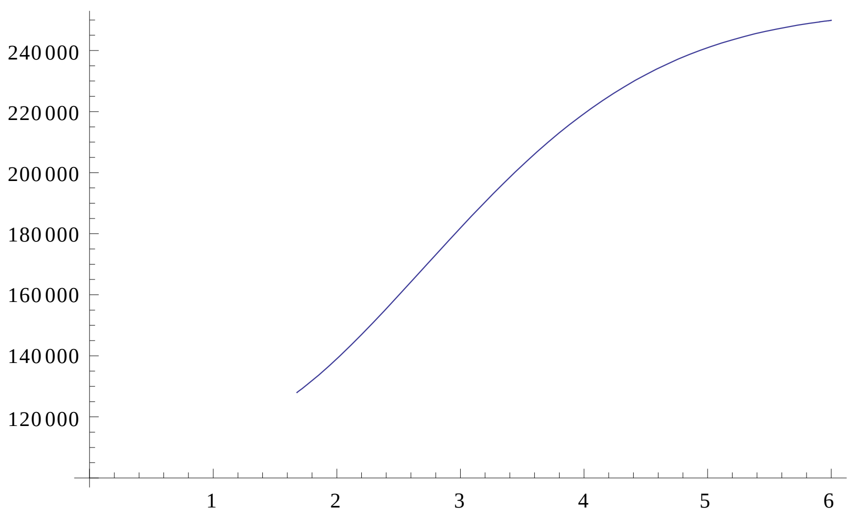

Figure 6. The bidding function, showing the optimal bid for a private signal of x.
(Do not bid if x < 1.67814.)

Note that the function tends to level off for extremely high signals. This is what we should expect. An evaluator with a private signal of 6 or more, who is willing to bid at least $240 thousand, has little risk of losing the auction, so it would take a very large (and very unlikely) increase in x to justify bidding higher.

i. Determining the outputs. After calculating the equilibrium strategy, it becomes possible to compute the outputs of the model.

We have already seen, in the example, how to compute the expected number of bids, and the probability of getting at least one bid.

Each of these quantities depends on the actual value of U, so we compute them conditional on U, and then integrate with respect to the distribution H. Given U=u, the probability of an evaluator getting a signal of at least x* (and therefore placing a bid) is

$$1 - F(x^* - u).$$

The mean number of bidders is m, and each bids with this probability, so the expected number of bids is m (1-F(x*-u)). Integrating over u gives

$$\text{Expected number of bids} = \int_u m\left(1 - F(x^* - u)\right) dH(u).$$

Similarly, the probability of getting a bid is

$$\text{Probability of at least one bid} = \int_u \left[1 - \exp\left(-m\left(1 - F(x^* - u)\right)\right)\right] dH(u)$$

The expected value of the highest bid, conditional on there being at least one bid, is

$$\text{Conditional high bid} = \int_{x \geq x^*} g^{-1}(x)\, p(x)\, dx,$$

where p(x) is related to the distribution of the highest bid, and is given by

$$p(x) = \int_u m e^{-m(1-F(x-u))} dH(u).$$

Empirical Testing

The bidding model with various assumptions about the signal generation process can be tested empirically. In the past, the Department has performed and commissioned a number of econometric studies of the determinants of competition and amounts bid in offshore lease sales, including [3] which includes a literature review in its Chapter 2.

The Interior Department maintains an extensive historical database of offshore leases bid on, number and dollar amounts bid, bidder identities, and physical characteristics of blocks. The bidding model presented in this paper can be tested with these data, with a focus on alternative signal generation models and other topics.

References

[1] IMODEL Technical Document, Version 5.0 (2000). Available from the US Minerals Management Service, Herndon, VA 20171.

[2] Richard B. Thompson and A. Larry Wright, "Equilibrium Bidding Strategies in Common-Value Sealed-Bid Auctions" (2004). Paper available from University of Arizona at math.arizona.edu.

[3] Ashton, P.K., L. O. Upton III, and Michael H. Rothkopf, “Effects of Royalty Incentives for Gulf of Mexico Leases,” 2 vols. (2005). US Interior Department, Minerals Management Service, OCS Study MMS-2004-077.

Appendix

Formal specification of the model

1. Rules of the Auction

The model applies to a first-price, sealed-bid auction. Conditions may include a bidding fee and a minimum bid. More precisely: Any party may submit a bid. Bidders do not know what other parties are considering bids or the values of other bids. The property is delivered to the party that submits the highest bid, provided that it is at least equal to the published minimum bid b_{min}. Ties being broken at random. If no bids of at least b_{min} are submitted, the property is not delivered to anyone.

We use the formal term "evaluator" for a potential bidder, since some of these parties may choose not to bid. The profit for an evaluator is given by

0	if the evaluator does not bid
–fee	if the evaluator bids but does not win
V – b – fee	if the evaluator bids $b \geq b_{min}$ and wins the auction.

Here V is the value of the property, described below. Each evaluator seeks to maximize the expected value of its own profit.[1]

We usually think of b_{min} and fee as non-negative, but we make no such formal requirement. We allow the value $b_{min} = -\infty$, meaning that all bids are considered. As there is no provision in current law for a bidding fee, we have assumed in most applications that fee = 0.

2. Common value assumption

We assume that the property has the same value V to all evaluators, but that V is unknown to them. We model V as a random variable subject to a prior distribution, representing all publically available knowledge. The participants use this distribution as a common prior

We would like to assume that the participants' estimating errors are independent of the value of V ("linear errors") but that is usually an implausible assumption. It is more plausible

[1] We are making a general assumption of risk neutrality on the part of the evaluators. This assumption is more palatable if all of our probability distributions are taken to be "martingale" or "risk-neutral" probabilities; that is, if they are inferred from market behavior rather than from repeated trials. We will not mention this subject again.

that some transformation of V that is subject to linear errors. We therefore introduce a random variable U called the "value parameter," and a function relating V to U:

$$V = v(U),$$

where v is increasing but not necessarily strictly increasing; that is, $v(u_2) \geq v(u_1)$ whenever $u_2 \geq u_1$. The function v is common knowledge, as is the cumulative distribution function H, defined by

$$H(u) = \Pr(U \leq u).$$

This representation allows for the distribution of U to be continuous or discrete, or a mix. Expected values dependent on H will be written as $\int \ldots dH(u)$. Readers who prefer a different notation may assume a density function h, and take "dH(u)" as a synonym for "h(u)du."

The distribution for U effectively defines a distribution for V.

3. Private information

Certain evaluators (selected by a process described below) receive private signals, and so become "informed evaluators." We denote the signal of the i-th informed evaluator by X_i, related to U by

$$X_i = U + R_i,$$

where R_i may be characterized as an estimation error. We model all of the X_i's and R_i's as random variables, and assume that the R_i's are independent of each other and of U --- that is, we are assuming independent linear estimating errors for U. The i-th evaluator knows the value of X_i, but not of R_i or any other evaluator's signal.

We assume further that estimation errors are normally distributed with a common standard deviation $\sigma > 0$ that is itself common knowledge. We write

$$c = 1/\sigma^2,$$

and refer to c (the inverse variance) as the "information parameter." The density function for each R_i is given by

$$f(r) = \sqrt{\frac{c}{2\pi}}\, e^{-cr^2/2}$$

(the normal density function with mean 0 and standard deviation σ). We denote the corresponding cumulative distribution function by F.

This is the same information model that Thompson and Wright use, except that we have specialized to the normal distribution for estimating errors (they allow general forms for f) and we apply the errors to U instead of to V directly. For much of what follows we will write f and F as though we allow general forms, but in the end we allow only normal distribution given above.

Normal distributions are plausible for estimating errors because that is what would typically result from a participant receiving many small signals, all subject to independent linear errors, and combining them. It is plausible to think of c as an “amount of information” available to the evaluators, since if one were to combine independent signals with different values of c, the effect would be to add the values of c. (We do not allow multiple signals for the same evaluator.)

4. The role of uninformed evaluators

Evaluators who do not receive private signals may still bid. They cannot expect to make a profit (on average) and under most conditions they choose not to bid and are irrelevant to the result. Sometimes, though, they do play a role.

Suppose, for example, that the lowest possible value for V, say v_0, is lower than the minimum bid b_{min}. Then it is rational for uninformed evaluators to bid v_0 (or slightly less) in the hope that their bid might squeak by. This forces informed evaluators to bid at least v_0; in effect, v_0 replaces b_{min} as the effective minimum bid.

We will calculate the value of b*, the optimal bid (if any) by uninformed evaluators, as part of our equilibrium analysis. When it matters at all, it replaces b_{min} in the initial condition used to compute the equilibrium strategy for informed evaluators.

5. How many evaluators? The information cost

We assume that the number N of informed evaluators is a random variable. Its distribution is common knowledge, and is given by

$$q(n) = \Pr (N = n)$$

for n = 0, 1, 2, …. We assume specifically that N is has a Poisson distribution

$$q(n) = (m^n/n!)\, e^{-m},$$

where the mean m is common knowledge. Note that there is a positive probability that N = 0; that is, that there are no informed evaluators. (This sometimes means that there are no bids, and sometimes means that the property goes to an uninformed evaluator.)

The user may specify m as a model parameter. We can go further, though, and internalize the determination of m as follows. First, we assume an “information cost” equal to k, also common knowledge.

Suppose that there are a very large number N of potential evaluators, each free to pay k and receive a private signal. The other parameters of the model being common knowledge, each evaluator can calculate the expected profit of an informed evaluator (as a function of m). At equilibrium, each potential evaluator will purchase information with some probability p chosen so that the profit of informed evaluators is exactly equal to k. In the limit of large N, the values of N and p do not matter except that they determine m = Np. The actual number of informed evaluators is then Poisson with mean m.

As the model is implemented, the user can either specify m (the mean number of informed evaluators) or k (the information cost). In the latter case, the mean m is determined within the model. (Essentially, we iterate the entire model until the expected profit of informed evaluators is equal to k.) The latter option has the advantage that the number of bids increases automatically for more valuable tracts. (That is, if H concentrates probability on high values of U, then m adjusts upward automatically.) This is what we would normally expect in auctions, and it is what we observe.

6. Summary of the Inputs

We summarize the inputs to the model:

b_{min}	minimum legal bid; may be $-\infty$
fee	bidding fee
H(u)	cumulative distribution parameter for the value parameter, U
v(u)	function giving value, V, in terms of the value parameter
σ	standard deviation of estimating errors
c	information parameter; same as $1/\sigma^2$
f(r)	density function for estimating errors, normal with mean 0, standard deviation σ
F(r)	cumulative distribution function corresponding to f(r)
k	information cost; must be positive
m	mean number of informed evaluators (specified by user only if k is not specified); must be positive
q(n)	probability that there are n informed evaluators; Poisson with mean m

All of these are assumed to be common knowledge. The parameter c and the functions f( ) and F( ) are determined by the choice of σ, and the function q( ) is determined by either k or m (whichever one the user chooses to specify).

In some of the formulas that follow we will use f, F, and q as if we are allowing general forms for these functions. Actually, as noted, we have defined specific forms for each of them.

7. The symmetric equilibrium

Under these conditions we derive a unique symmetric equilibrium strategy for the evaluators. The outputs from the model arise as consequences of the general equilibrium.

We do not consider the possibility of an asymmetrical equilibrium.

A symmetric equilibrium is defined by two (or three) numbers and a function:

m	Mean number of informed evaluators (if not specified by user)
b*	The highest rational bid (if any) by an uninformed evaluator
x*	The lowest private signal that justifies a bid
g(x)	The optimal bid by an informed evaluator with private signal x.

If the uninformed evaluators do not bid, we take $b^* = b_{min}$. The function g(x) is then defined for all $x \geq x^*$ and satisfies $g(x^*) = b^*$. We assume that g(x) is strictly increasing over this range; in fact, we assume that the derivative g'(x) exists and is positive for $x > x^*$.[2]

(Under the assumptions of the model the informed evaluators do not need to consider mixed strategies. Neither do the uninformed evaluators, if the bidding fee is zero. When there is a positive bidding fee, however, it may be rational for uninformed evaluators to bid randomly within a certain range. Clarifying this situation would require a more precise model for this special case. It isn't worth the trouble. In effect we assume in this case that there is just one uninformed evaluator, and with this simplification, mixed strategies are unnecessary and b* is well defined.)

Solution of the model requires searches for the equilibrium values of m, b*, and x*. Given values of these numbers, the function g( ) is determined by numerical solution of a differential equation.

In the actual procedure, we first guess values of m and b*. (Initially we choose $b^*=b_{min}$; usually this is the final value as well.) We compute x* by a one-dimensional search, and then compute g( ) in tabular form by numerical solution of a differential equation. Having g( ), we compute a better value for b*, and iterate until b* converges. We then compute the expected profit of informed evaluators. We repeat the entire process for different values of m until that expected profit is equal to the information cost k. In practice the searches are not difficult and they converge quickly.

[2] We derive g(x) using standard methods, and the derivation constitutes a proof that any symmetric equilibrium must be characterized by the b*, x*, and g( ) that we obtain. We do not prove, however, that g(x) always has a positive derivative. It has this property in every case we have tried, and in every case we have tried using similar models elsewhere. We cannot rule out the possibility that for some combination of inputs, the derived g( ) would fail to be increasing. In that case we would throw up our hands and declare that the model had failed.

Here are the formulas for the symmetrical equilibrium strategy, given the assumptions of the model.

Given m, and b*, the expected profit for an informed evaluator with private signal x and a bid of b is

$$\pi(b) = -\text{fee} + \frac{\int_u \sum_n nq(n)\left[v(u)-b\right] f(x-u)F(x-u)^{n-1}dH(u)}{\int_u \sum_n nq(n) f(x-u)dH(u)}.$$

For a fixed value of b, this expression normally increases as a function of x. We choose x* to be the smallest value of x for which π(b*) ≥ 0; that is, it is the smallest private signal that justifies a minimum bid.

Now given m, b*, and x*, the function g( ) is the solution to the differential equation

g'(x) = c(x) – a(x) g(x) (for x ≥ x*)

subject to the initial condition

g(x*) = b*.

(We check that the resulting function g is strictly increasing; if not, then the model has failed.) The coefficient functions a and c are given by

$$c(x) = \frac{\int_u \sum_n n(n-1)q(n)v(u) f(x-u)^2 F(x-u)^{n-2}dH(u)}{\int_u \sum_n nq(n) f(x-u)F(x-u)^{n-1}dH(u)}$$

$$a(x) = \frac{\int_u \sum_n n(n-1)q(n) f(x-u)^2 F(x-u)^{n-2}dH(u)}{\int_u \sum_n nq(n) f(x-u)F(x-u)^{n-1}dH(u)}$$

(These expressions are not nearly as complicated as they look! One can interpret c(x)/a(x) as the expected value of the property, given that this evaluator has private signal x and that it is tied for highest signal.)

Knowing m, x*, and the function g, the profit of an uninformed evaluator bidding b is given by

$$\pi_u(b) = -\text{fee} + \int_u \sum_n q(n)\left[v(u)-b\right] F(g^{-1}(b)-u)^n dH(u).$$

We determine b* to be the value of b that maximizes this function, provided that it is at least equal to b_{min} and makes the profit positive. (Otherwise we stop with $b^* = b_{min}$.) We repeat this process until b* converges.

Finally, knowing m, x*, b*, and the function g, we can compute the expected profit for informed evaluators. We repeat the entire process for varying values of m until this expected profit is equal to k.

8. Outputs from the model

The model is rich enough to determine a probability distribution over all possible outcomes. In practice, we compute three values:

The expected number of bids;

The probability that there is at least one bid; and

The expected value of the highest bid, conditional on there being a bid.

The expected revenue to the seller is the product of the last two values.

In the absence of a bid from an uninformed evaluator, the outputs are given by the following formulas.

The expected number of bids is

$$\text{ENB} = \int_u m\left(1 - F(x^* - u)\right) dH(u).$$

The probability that there is at least one bid is

$$\text{PBID} = \int_u \left[1 - \exp\left(-m\left(1 - F(x^* - u)\right)\right)\right] dH(u).$$

Conditional on there being at least one bid, the expected value of the highest bid is

$$\text{CHB} = \int_{x \geq x^*} g^{-1}(x)\, p(x)\, dx,$$

where p(x) is related to the distribution of the highest bid, and is given by

$$p(x) = \int_u m e^{-m(1-F(x-u))} dH(u).$$