

NBER WORKING PAPER SERIES

ENTRY AND SELECTION IN AUCTIONS

James W. Roberts
Andrew Sweeting

Working Paper 16650
<http://www.nber.org/papers/w16650>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2010

We are extremely grateful to Susan Athey, Jonathan Levin and Enrique Seira for sharing the data with us and to Coleman Doggett and Doug MacDonald for their insight on the timber industry. We appreciate useful feedback from John Asker, Phil Haile, Jakub Kastl, Brian McManus, Rob Porter, Katja Seim and numerous seminar participants as well as attendants of the 2010 NBER Summer IO program. Excellent research assistance was provided by Christoph Bauner, Yair Taylor and Huihui Wang. Both authors gratefully acknowledge financial support from Duke University. Any errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by James W. Roberts and Andrew Sweeting. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Entry and Selection in Auctions
James W. Roberts and Andrew Sweeting
NBER Working Paper No. 16650
December 2010
JEL No. D44,L20,L44,L73

ABSTRACT

We develop and estimate an entry model for second price and open outcry independent private value auctions where potential bidders receive an imperfectly informative signal about their value prior to deciding whether to pay a sunk entry cost. In this way the model flexibly allows for selection on values, which will affect an entrant's subsequent competitiveness, at the entry stage. As signals become more informative, the entry process exhibits greater selection as firms with higher values are more likely to enter. We allow for asymmetries across bidders and unobserved heterogeneity across auctions, which are important features of most data sets. We show how incorrectly assuming the extremes of either no selection (no signal) or perfect selection (prior knowledge of one's value) - the common alternatives in the literature - can yield incorrect estimates of model primitives and bias the results from counterfactuals. We apply our model to U.S. Forest Service timber auctions and find strong evidence in favor of a selective entry process. We take advantage of the flexible entry model to reevaluate the well known theory result that with fixed participation a seller prefers an extra bidder over the ability to set an optimal reserve price. In our model, the relative value of setting a reserve price and increasing the number of potential entrants to a revenue-maximizing seller will depend on the degree of selection. Our structural estimates imply that, if the USFS wants to maximize revenues, it will benefit more from adding an additional potential entrant than setting an optimal reserve price.

James W. Roberts
Department of Economics
213 Social Sciences
Box 90097
Duke University
Durham, NC 27708
j.roberts@duke.edu

Andrew Sweeting
Department of Economics
213 Social Sciences
Box 90097
Duke University
Durham, NC 27708
and NBER
atsweet@duke.edu

1 Introduction

Empirical entry models typically ignore the possibility that the entry process tends to select entrants with unobserved characteristics, such as lower marginal costs or higher product quality, that make them subsequently more competitive.¹ As a result the possibility that firms choose to participate because they know that they are likely to be particularly competitive is ruled out. In this paper, we set-up an estimable selective entry model for an empirically important setting - auctions - and show that allowing for selective entry can have important implications for the estimation of the distribution of bidder values and counterfactuals. We estimate our model using data from U.S. Forest Service (USFS) timber auctions in California where we find strong evidence in favor of selection. We use our estimates to reevaluate, in the context of entry, a well known result from auction theory (see Bulow and Klemperer (1996)) that with fixed participation a seller prefers an extra bidder over the ability to set an optimal reserve price. In our model, the relative value of setting a reserve price and increasing the number of potential entrants to a revenue-maximizing seller will depend on the degree of selection. Our structural estimates imply that, if the USFS wants to maximize revenues, it will benefit more from adding an additional potential entrant than setting an optimal reserve price.

In many auction settings bidders must make sunk investments to better understand the potential returns to participating in the auction. A few examples (which correspond to the leading applications in the empirical auction literature) include timber auctions where bidders need to “cruise” the tract they are bidding on to learn about the trees they will harvest, procurement auctions, such as those for highway paving contracts, where bidders must assess the project in order to gauge the likely cost of completing the job and offshore oil tract auctions where bidders conduct seismic surveys to form expectations about the amount of oil present. Despite their importance, entry processes have only recently been incorporated into empirical auction research. When models include these, the assumptions made about bidders’ information are often extreme. For example, one assumption is that bidders have no private information about their value prior to paying the entry cost. An alternative assumption is that bidders know their value perfectly before paying an entry cost. These models are based on Levin and Smith (1994) and Samuelson (1985), respectively. Throughout the paper we refer to these as the LS and S models. These polar cases are rarely going to be correct and a more reasonable modeling assumption lies somewhere in between. This is the type of model we explore in this paper.

Specifically, we estimate a two-stage entry model for second price or ascending independent private value (IPV) auctions with asymmetric bidder types which we refer to as the Signal model.² In the first stage potential bidders simultaneously decide whether to participate in the auction,

¹The typical formulation in the static entry literature (e.g., Berry (1992)) allows that firms with higher ε draws are more likely to enter. However, these ε draws do not directly affect the profitability of other firms, and so they are more appropriately thought of as shocks to fixed costs or sunk entry costs that are irrelevant to market outcomes conditional on entry. Similarly, the dynamic entry literature assumes that entrants receive iid shocks to their entry costs but are otherwise symmetric.

²The focus on a second price model makes the exposition easier. Our empirical work allows for the fact that the auctions we use are open outcry auctions rather than second price sealed bid auctions.

which entails paying a sunk entry cost that enables them to learn their value exactly. In the second stage, entrants submit bids, which in the equilibrium of a second price auction will be equal to their values. We allow for selection by assuming that each potential entrant gets a private information signal about its value before the entry decision is made. Equilibrium for each type of bidder is characterized by an entry threshold whereby their signal must be sufficiently optimistic about their value to justify paying the sunk entry cost. The signal's precision determines how selective (on values) the entry process is. If the signal is very uninformative, then the marginal and inframarginal value distributions will be very similar (little selection) and outcomes will approach those of the LS model. On the other hand, when signals are very informative, these distributions will be quite different (a great deal of selection) and outcomes will approach those of the S model.

We believe that this set-up corresponds to how one might think about entry decisions in both timber and other auction settings. In the timber auction setting, bidders have different values for timber tracts reflecting their individual capacity utilization, specialities and downstream contracts. The potential entrants' own knowledge of the forest, together with the information provided by the government prior to the auction (e.g., its estimates of wood volume and species type), will allow a potential bidder to form an estimate of his value (the signal). If it decides to enter, then the firm will do its own survey of the tract and undertake other costly investments (which represent the entry cost) to exactly ascertain its value.

The degree of selection matters because it determines the difference between the values of the marginal and infra-marginal entrants. Ignoring this difference will bias estimates of the distribution of values, which is the central object of interest in much of the empirical auction literature. For example, ignoring selection may lead one to believe that values are higher and less dispersed than they really are because the set of values that are observed tend to be drawn from the upper part of the distribution. In a setting with asymmetric bidder types, one will tend to overestimate the mean values of weaker types because these firms are particularly likely to enter when they get especially high signals. This may mask differences across bidder types and also affect the predictions of any counterfactual that impacts the entry margin.

For example, consider a subsidy to encourage greater participation in an auction. With selection, the marginal entrant who is attracted into the auction by the subsidy is less likely to be valuable to the seller than in a model without selection where the marginal and inframarginal entrants look alike. The degree of selection also affects what happens to sellers' revenues as the number of potential entrants increases. As shown in the original LS paper with symmetric firms, expected revenues decrease with the number of potential entrants in a model with no selection once firms enter with probability less than one.³ In contrast, revenues tend to increase when there is enough selection. We use this difference to support our contention that there is a selective entry process in our data. Another reason that selection may matter is that it impacts the optimal reserve price, a crucial tool in mechanism design. The well known result that, with a fixed set of bidders, the seller's optimal reserve price is independent of the number of bidders and should be set above his

³The symmetric equilibrium in a symmetric LS model involves firms using mixed strategies.

value (see Myerson (1981) and Riley and Samuelson (1981)) breaks down when entry is introduced. Moreover, the new optimal reserve pricing rule is not robust to different models of entry. In the LS model it is optimal to set the reserve price independent of the number of bidders and at the seller’s value, while in the S model it is optimal to set it above the seller’s value and increase it in the number of bidders. For this reason it is prudent to be flexible in the amount of selection a model permits.

In our empirical application we estimate our model using several approaches (Nested Fixed Point, Nested Pseudo-Likelihood and Simulated Maximum Likelihood with Importance Sampling) that differ in the extent to which they allow us to handle asymmetries in the parameters across bidder types (in our example, timber mills and logging companies), heterogeneity across auctions and the possible existence of multiple equilibria. To our knowledge the nested pseudo-likelihood procedure has yet to be applied to an auction environment. Therefore, through a variety of Monte Carlo experiments, we show that this latter estimator performs well in uncovering the true parameters of the data generating process, even when it approximates the LS or S models. Consistent with reduced form evidence, the coefficient estimates suggest quite a selective entry process, resulting in a marked difference between the value distributions of the marginal and inframarginal entrant.

We take advantage of the flexible entry model to reevaluate the well known theory result regarding the benefits of competition. Specifically, Bulow and Klemperer (1996) show that with fixed participation a seller prefers an auction with no reserve to one with an optimal reserve but one less bidder. That is, the seller is willing to give up his main revenue extraction tool and potentially “give away his object for nothing” in favor of attracting one more bidder. Extending this result to settings with costly entry is not straightforward. In particular, depending on the amount of selection in the entry process, the seller may not prefer to increase the number of potential entrants if he must lose the ability to set an optimal reserve price: increasing competition does not necessarily justify the seller relinquishing all of his bargaining power. Our flexible model of entry, therefore, can be usefully employed to understand the benefits of competition when entry is costly. Our structural estimates imply that, if the USFS wants to maximize revenues, it will benefit more from adding an additional potential entrant than setting an optimal reserve price.

The paper proceeds as follows. Section 2 discusses the relevant literature, Section 3 introduces our model and discusses its identification, Section 4 illustrates the importance of selection, Section 5 presents our estimation methods, Section 6 turns to our empirical application, Section 7 uses our estimates to compare the benefits of auctions and negotiations to the USFS and Section 8 concludes. The appendices include alternate estimation methods and results.

2 Literature Review

In the empirical auction literature (see Hendricks and Porter (2007) for a recent survey) bidders’ decisions of whether or not to participate in an auction are frequently ignored. When the set of bidders is not considered fixed, the most common entry model used is the LS model. Examples

include Athey, Levin, and Seira (forthcoming) who examine timber auctions, Bajari and Hortaçsu (2003) who consider coin auctions on eBay and Krasnokutskaya and Seim (forthcoming) who analyze procurement auctions. Palfrey and Pevnitskaya (2008) and Ertuç, Hortaçsu, and Roberts (forthcoming) analyze entry in auctions using experimental data and some structural techniques. In a recent paper Li and Zheng (2009) compare the effects on competition in procurement auctions stemming from both the LS and S models. On the other hand, we are interested in a less restrictive model that does not assume these extreme entry models. Such a model has been proposed in the literature before. For example, Hendricks, Pinkse, and Porter (2003) consider such a model in their testing of competitive equilibrium bidding in offshore oil auctions. In a recent paper Marmer, Shneyerov, and Xu (2010) consider testing whether the LS, S or a general affiliated signal model of entry best explains bidding behavior in procurement auctions. The idea behind their test is, relying on exogenous variation in the number of potential bidders, to examine whether the distributions of entrants' valuations varies with the number of potential bidders. Using the same data as Li and Zheng (2009), they find support for the S model and their affiliated signal model. They also estimate a very simple version of their model with symmetric bidders and no unobserved auction heterogeneity, a feature that has been shown to be important in other settings (Krasnokutskaya (forthcoming), Athey, Levin, and Seira (forthcoming), Roberts (2009), Hu, McAdams, and Shum (2009)). In this paper we estimate a fully structural model of entry and bidding in independent private value ascending auctions with asymmetric bidders and unobserved heterogeneity. The objects of interest will be bidders' true value distributions (that is the value distribution of any potential bidder, not only a participant), signal distributions and entry costs. The aim is to illustrate the biases in demand estimation and counterfactual analysis when the incorrect model of entry is employed.

One of our goals is to effectively incorporate bidder asymmetries into the model and estimation. With asymmetric bidders, the choice of a flexible entry model may be particularly important. For example, consider an LS model with asymmetric types under the assumption that all firms within a type use the same strategy (type-symmetric). When there are multiple types, no type-symmetric equilibrium in this model can have some, but not all, firms entering within more than one type. This might make it impossible to rationalize what is observed in many data sets without introducing additional shocks to entry costs. A Signal model can explain these patterns by firms receiving different signals about their values.

Einav and Esponda (2008) and Bhattacharya, Roberts, and Sweeting (2010) estimate this paper's imperfectly selective entry model when the post entry game is a first price auction. This case presents the increased computational challenge of solving the equilibrium bid functions along with the entry decisions. Both papers adapt a parametric approach and the latter allows for unobserved auction heterogeneity.

Performing counterfactual analyses is one goal of estimating our structural model. In related work, Brannman and Froeb (2000) use a sample of 51 Oregon timber auctions to evaluate various counterfactuals such as mergers and bidder preference programs. The entry margin, which they

ignore, is likely to be particularly important for programs that encourage entry by weaker bidders such as bid subsidies or preference programs. Krasnokutskaya and Seim (forthcoming) apply the LS model to highway procurement to assess current bidder preferential treatment programs where small contractors receive subsidies to make them more competitive with their larger rivals. Their analysis finds that the effects of the program and its resulting impact on government costs vary by project, but in general, smaller bidders have an increased chance of winning. Assuming the opposite model of entry, Hubbard and Paarsch (2009) simulate a theoretical model in order to evaluate (computationally) bid preference programs when bidders know their value prior to participation. They find that the effect of preference programs on bidder participation is relatively unimportant.

This paper is also related to the literature on estimating incomplete information entry games. In the canonical entry model in this literature, each firm receives a private information shock to its entry cost, and in equilibrium firms choose to enter if their draw is high enough. In our model, potential bidders choose to enter if their signal about their value is high enough. However, an important difference is that in our model the signal and the entry decision are correlated with how competitive the firm is once it enters, whereas in the standard model the draw on the entry cost has no effect on how competitive the firm is once it enters. Instead, the natural analogue of our model in other entry settings would be one where firms receive noisy signals about their post-entry marginal costs or qualities.⁴ Roberts and Sweeting (2010a) adapts the framework presented in this paper to a non-auction setting: airlines. There, the goal is to measure the true effect that potential entrants have on limiting prices in a Cournot game when actual market participants are a selected set of firms.

Finally, in our empirical application we focus on auctions for the right to log federally owned forestland. There is now a long line of empirical auction literature analyzing these auctions (see Paarsch (1997), Baldwin, Marshall, and Richard (1997), Haile (2001) and Athey and Levin (2001) to name a few). We are the first to employ our more general model of entry to these data. In a companion paper (Roberts and Sweeting (2010b)) we evaluate the USFS's returns to switching their sales mechanism from the current format in which bidders compete simultaneously to one in which the USFS approaches bidders sequentially. This alternate format allows bidders to preempt future potential entrants from participating by submitting higher bids earlier in the process. We find that when entry is selective, as we indeed find that it is in these timber auctions, it is often the case that the USFS could improve revenues and efficiency by switching their allocation mechanism to a sequential format.

⁴Of course, one might consider a model that combined both an imperfectly informative signal and a random shock to entry costs. However, if both sets of draws are iid then it is not obvious that the precision of the signal and the variance of the entry cost shocks can be separately identified except by functional form assumptions. We discuss identification of our model in Section 5.2.

3 Model

We now present our model of entry into auctions. We begin by introducing the model and showing that it is characterized by a cutoff strategy whereby bidders only enter auctions when their signal is sufficiently high. For much of the discussion we refer to the mechanism as being a second price auction. However, given our informational assumption that bidders have independent private values, the strategies introduced here are akin to those in an English Button auction.⁵

3.1 A General Entry Model with Selection

We consider a series of $t = 1, \dots, T$ second price IPV auctions. In any auction there is a set of potential bidders who may be one of $\tau = 1, \dots, \bar{\tau}$ types. Let the number of potential bidders of any type τ be N_τ and the number that choose to enter the auction be n_τ . Bidder values V are distributed according to $F_\tau^V(V)$ where the dependence on their type is made explicit. For any auction a the object may differ according to observable and unobservable dimensions and thus we allow the distribution to be auction specific $F_{\tau a}^V(V)$. One observable dimension along which auctions may differ is the seller's reserve price. Bidders know the distributions from which their own and their competitors' values are drawn. There are two stages to the game, an entry stage and an auction stage.

At the start of the entry stage every bidder observes the set of potential bidders and a signal S (we drop the i subscript for now) which is affiliated with their value of the object. If they pay an entry cost K_τ they learn their value for sure. For expositional ease, we will focus on the special case where $S = V + \varepsilon$, $V \sim N(\mu_\tau, \sigma_V^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. In practice σ_ε^2 and σ_V^2 may also depend on τ . Therefore $E[S|V] = V$, $\varepsilon \sim F^\varepsilon(\cdot)$ and $E[\varepsilon|V] = E[\varepsilon] = 0$.⁶ Given these assumptions, we can characterize a bidder's posterior value distribution after observing a signal s . In particular, the bidder now believes his value is drawn from $N(\alpha\mu_\tau + (1 - \alpha)s, \sigma_V'^2)$, where $\alpha = \frac{1/\sigma_V^2}{1/\sigma_V^2 + 1/\sigma_\varepsilon^2}$ and $\sigma_V' = \sqrt{\frac{1}{1/\sigma_V^2 + 1/\sigma_\varepsilon^2}}$. During this entry stage all bidders simultaneously decide whether or not to pay a fixed cost K_τ to observe their true value for the object. Any bidder that does pay K_τ proceeds to the auction stage of the game.⁷ In some simulations and in our empirical example we consider an alternative model where $S = VA$, $A = e^\varepsilon$, $V \sim \log N(\mu_V, \sigma_V^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$.

At the auction stage of the game bidders who paid K to learn their true value participate in the auction. In our second price auction setting, bidders submit their true value regardless of how many bidders entered. The object is awarded to the highest bidder at a price equal to the second highest bid or the reserve price (which bidders knew at the entry stage), whichever is greater.

A bidder's strategy in this game must include a rule for entering, as a function of his signal, and a

⁵When we introduce our estimation strategy in Section 5, we extend our methodology to cover more general models of bidding in open outcry auctions which won't require all bidders to bid up to their true value. We are also currently working on extending the analysis to first price auctions.

⁶Many of our results, and indeed the intuition for our modeling contribution, could be established with a more general model of the relationship between a bidder's signal and his value (perhaps they are affiliated, for example). However, since our ultimate goal is to take the model to the data, for clarity sake we present this more specific model.

⁷Here we assume that any bidder must pay K_τ to participate in the auction.

rule for bidding as a function of his value and (potentially) his signal. In a second price auction, the dominant bidding strategy is to bid one's value. An optimal entry strategy must involve entering if and only if the signal lies above some threshold, as shown in the following proposition. This threshold signal is implicitly defined by the zero profit condition whereby the bidder receiving it is indifferent between staying out and paying the entry cost and entering the auction.⁸

Proposition 1. *The optimal bidding strategy for entrants is to bid one's value. The optimal entry strategy is given by a signal threshold s'_τ for any type τ bidder such that he pays the entry cost if and only if his signal $s > s'_\tau$.*

Proof. Let $\mathcal{M}_{-i}$ be the set of entrants other than bidder i . Entrants submit a bid to maximize their expected profit conditional on entry:

$$(v_i - E[\max\{v_{-i}, r\}] : v_j \leq b, \forall j \in \mathcal{M}_{-i}) \prod_{\mathcal{M}_{-i}} F(b|\tau)$$

Standard arguments show that participants have the dominant strategy to bid their value regardless of the number of potential bidders or asymmetries among them. This strategy is equivalent in an ascending auction. Profits are increasing in a bidder's value. Because signals and values are affiliated, a higher signal leads the potential entrant to raise his beliefs about his value and because signals are independent across bidders, it does not alter his beliefs about other bidders' post entry competitiveness. Thus, for any signal at which the bidder enters, he would enter for any higher signal and for any signal at which he does not enter, he would not enter for any lower signal. Therefore an equilibrium entry rule follows the threshold rule. An equilibrium exists because any bidder's reaction function is continuous in his and his opponents' thresholds. *Q.E.D.*

If $\bar{\tau} = 1$, then there is one symmetric equilibrium, which is also true in the LS and S models. This is established in the following proposition.

Proposition 2. *With one type of potential entrant there will be a unique symmetric equilibrium.*

Proof. Suppose that there are two symmetric equilibria, where potential entrants have cutoffs s_1^* and s_2^* respectively and $s_1^* > s_2^*$. Consider a potential entrant i . For any v_i , i 's expected profits from entering will be increasing in s_{-i}^* , the cut-off used by all other players as for any set of v_{-i} s an increase in s_{-i} reduces the probability that rivals will enter. From this it follows that i 's best response cutoff to s_{-i}^* is decreasing in s_{-i}^* . If so, if s_1^* is i 's best response to s_1^* , it cannot also be the case that s_2^* is i 's best response to s_2^* , so that s_1^* and s_2^* cannot both be symmetric equilibrium thresholds. *Q.E.D.*

While the above proposition proves the uniqueness of a symmetric equilibrium with one bidder type, there can still be asymmetric equilibria. However, only the symmetric equilibrium exists as the number of potential bidders grows large.

⁸We explicitly define the zero profit condition in Section 5.

3.1.1 Multiple Equilibria with Asymmetric Bidders

Entry models typically have multiple equilibria that differ according to which firms enter. Even when we focus on “type-symmetric” equilibria (i.e., ones in which all firms of a particular type using the same cut-off) and second price auctions (with a dominant strategy equilibrium in the bidding game), there can be multiple equilibria for some values of the parameters. The possible existence of multiple equilibria influences our empirical specifications and our choice of estimation techniques.

To see why multiple equilibria arise, consider a simple case where there are two types of firms ($\bar{\tau} = 2$), and one potential entrant of each type. Values are distributed $N(\mu_\tau, \sigma_V^2)$ for each type, with $\mu_1 = 200$ and $\sigma_V^2 = 25$. For both types, signal noise is normally distributed with $\sigma_\varepsilon^2 = 10$, and $K = 20$. The left panel of Figure 1 shows the best response functions for each firm, where $\mu_2 = 200$ so that the types are identical and the reserve price is 100. The s ’s where the reaction functions cross are Bayesian Nash equilibria and because the types are identical in this example, there is a symmetric equilibrium (on the dashed 45° line) where both types have the same strategies. However, there are also two asymmetric equilibria. At the top-left intersection of the reaction functions, the type 2 firm enters if it gets a signal above 115, which is almost certain for these parameters, whereas the type 1 firm enters only if it gets a signal greater than 218, which is less likely. The entry thresholds are reversed at the bottom-right intersection.




Figure 1: Reaction functions for symmetric and asymmetric bidders. In the left panel there are two symmetric bidders with values distributed $N(\mu, \sigma_V^2)$, with $\mu = 200$, $\sigma_V^2 = 25$, $\sigma_\varepsilon^2 = 10$, and $K = 20$. In the right panel we introduce asymmetries only in means; all other parameters are as in the left panel. The thin lines correspond to means of 200 and 190 and the bold lines correspond to means of 200 and 160.

The right panel of Figure 1 shows what happens when we introduce asymmetry among the firms. Here the two types have different mean values but share the other parameters. Type 2 is assumed to have a lower mean value, with either $\mu_2 = 190$ or $\mu_2 = 160$. A reduction in μ_2 causes type 2’s reaction function to move to the right. However, it also causes type 1’s reaction function to shift

down, i.e., it becomes more willing to enter, since for a given s'_2 type 2 is likely to have a lower value if it enters, which makes type 1 entry more profitable. Given the inverse S-shape of the reaction functions, which reflects the shape of the underlying value and error distributions, there remain a maximum of three possible equilibria. However, the direction of the shifts mean that there is only one equilibrium where the type 1 firm, with the higher mean value, has the lower s' . An equilibrium of this form continues to exist even when the difference in mean values is large enough that there is a unique equilibrium, which is the case when $\mu_2 = 160$.⁹ This type of equilibrium seems more plausible in the sense that the higher value types should hold the advantage, and it will tend to lead to more efficient allocations.¹⁰

When we adopt estimation methods that require an equilibrium selection model, such as Maximum Likelihood using a Nested Fixed Point algorithm and Simulated Maximum Likelihood using Importance Sampling, we will assume the equilibrium played is the one where the higher value types have the lower s 's and that the parameters apart from μ are common across the types (even if they differ across auctions) since otherwise there may be multiple equilibria where the high type has a lower s' .¹¹ However, this selection rule is an assumption. We therefore also consider an estimation technique - Nested Pseudo-Likelihood - that is (at least potentially) more robust to multiple equilibria in the sense that it may be able to converge to the equilibrium played in the data.¹² In doing so, we can also relax the assumptions that the parameters σ_V^2 , σ_ε^2 and K are common across the types, as we no longer need assumptions that will generate only one equilibrium of a particular kind.

3.1.2 Marginal and Inframarginal Bidders with Selection

We now illustrate how bidders' observation of a signal about their values "selects" the types that enter. We do this by examining the difference between the value distributions of entrants and those bidders who received a signal just low enough to stay out of the auction, what we term to be the "marginal bidder". This is an important distinction when considering the benefit of the marginal bidder participating, say due to a subsidy. In the LS model these distributions are the same. In the Signal model, however, the bidder receiving a signal $s = s'$ will have a value distribution that is first order stochastically dominated by the value distributions of entrants. This is clearly shown in

⁹Other changes in the parameters can also tend to lead to a unique equilibrium. For example, small entry costs, dispersed values or very uninformative signals can all tend to flatten the reaction functions leading to single intersection/equilibrium.

¹⁰Note that the analysis does not depend on having a single firm of each type. With multiple firms there will still be a symmetric equilibrium when the types have equal values and a decrease in the mean value of the second type will still cause the equilibrium best response function (i.e., the best response that solves only for the equilibrium among players of the same type) will shift to the right and the equilibrium best response function of the type 1 players will shift down. However, with multiple players the best response curves can tend to be flatter which also tends to favor the game having a unique equilibrium.

¹¹An alternative approach would be to follow Sweeting (2009) and Bajari, Hong, and Ryan (forthcoming) by estimating an equilibrium selection mechanism that allows for a mixture of different equilibria to be observed in the data.

¹²Pesendorfer and Schmidt-Dengler (2010) note that because Nested Pseudo-Likelihood is based on best response iterations it may not be robust if the equilibrium played in the data is not locally stable.

Figure 2. The left panel of Figure 2 displays the average value distributions of entrants and marginal bidders in the Signal with a precise signal and LS model. The right panel of Figure 2 displays the average value distributions of entrants and marginal bidders in the Signal with a less precise signal and LS model. It is clear that as the signal noise increases, the Signal model approaches the LS model.




Figure 2: Each figure compares value distributions for marginal bidders and entrants in the LS and Signal model. For each $V \sim N(200, 25)$ with $K = 20$ and $N = 6$. In the signal model $S = V + \varepsilon$, $\varepsilon \sim N(0, \sigma_\varepsilon)$. The left panel compares the value distributions for “marginal” bidders according to LS model and a Signal model with a lot of selection: $\sigma_\varepsilon = 5$. The right panel compares the value distributions for “marginal” bidders according to LS model and a Signal model with a lot of selection: $\sigma_\varepsilon = 50$.

4 Illustrating the Importance of Selection

In this section we illustrate why allowing for selection effects at the entry stage is important. To do this we first show how applying the incorrect entry model to the data causes the econometrician to incorrectly estimate demand. Second, we consider a relevant counterfactual question involving entry subsidies to show how this type of policy-relevant question may be incorrectly addressed when too restrictive models are taken to the data.

Incorrectly Estimating Demand

As stated above, far and away the most commonly used entry model in empirical auction work is the LS model. Here we consider what happens to parameter estimates when the true data generating process is the Signal model but the econometrician estimates the LS model. Table 1 shows the estimated parameters of the value distribution when the family is correctly chosen (that is a normal

N	$\hat{\mu} = 120$	$\hat{\sigma}_V = 25$	$\hat{K} = 10$	s'
3	132.0222	17.8262	20.8270	106.8737
4	137.3831	16.2657	18.2699	116.1196
5	141.0022	15.4720	17.5082	122.1247
6	143.6077	14.5615	17.0921	126.4939
7	146.2489	14.2386	16.8575	129.8908
8	147.8688	13.4926	16.2295	132.6493
9	149.4981	13.2966	15.7235	134.9597
10	150.9570	12.9902	15.3284	136.9395

Table 1: Incorrectly Estimating Demand, LS Model. The table displays estimates of model parameters assuming the LS model. Based on generating $T = 5000$ auctions where $S = V + \varepsilon$, $V \sim N(\mu, \sigma_V^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. Here $\sigma_\varepsilon = 5$. We solve for s' as the equilibrium outcome for this game for each value of N and this appears in the last column.

distribution is correctly assumed) but the LS is incorrectly applied. We generated data using our Signal model for given parameters (shown in the table) and then estimated μ_V , σ_V , K .¹³

The intuition for the results in Table 1 is straightforward. By employing the LS model, the econometrician wrongly assumes that bids observed in the data are placed by representative bidders since, supposedly, there is no selection effect. In reality, bidders with high signals chose to enter and this is correlated with higher values, thereby biasing the econometrician’s estimates of the value distribution. In particular, the estimated distribution under the LS model will first order stochastically dominate and be second order stochastically dominated by the distribution under the Signal model.¹⁴ To understand the consequences for estimating the entry cost, one needs to consider the two models’ predictions about entry. The LS model predicts greater equilibrium entry than does the Signal model since the greater dispersion of values of actual entrants increases the option value of participating for potential entrants. For example, at the true value distribution and entry cost parameters, if bidders have no signal about their value (the LS model), they enter 50.4% of the time. If they have a signal distributed as in Table 1, bidders enter 39.9% of the time. With an overestimated mean value, this effect will be compounded. There will also be a countervailing force due to the underestimate of σ_V . This is because the surplus conditional on winning is smaller the lower the standard deviation of values. On net, these three effects lead the LS model to predict more entry than is observed in the data and it justifies this “low level of entry” by overestimating K . This bias in K will fall as N increases since the LS model won’t predict as much of an increase in entry as does the Signal model and so the increased entry will be partially attributed to lower and lower entry costs.

We can consider a similar experiment with the S model. The results of incorrectly assuming the data generating process is the S model, when in fact it is the more general Signal model, appear in Table 2. The table shows the estimated parameters of the value distribution when the family

¹³Here we imagine that we have a very large data set over which to estimate the parameters. Analogously, we could set T to be smaller but replicate the small set of T many times.

¹⁴The second order stochastic dominance follows immediately from $\sigma_V > \sigma'_V$ whenever $\sigma_\varepsilon > 0$.

N	$\hat{\mu} = 120$	$\hat{\sigma}_V = 25$	$\hat{K} = 10$	$\hat{s}'$	s'
3	115.9863	29.7016	7.0017	97.5955	106.8737
4	113.1916	31.7532	7.4736	106.1968	116.1196
5	111.6691	32.5119	8.0142	112.9868	122.1247
6	109.6941	33.6069	8.2450	117.1782	126.4939
7	107.8251	34.5973	8.2982	120.2677	129.8908
8	106.2405	35.2924	8.3522	123.0219	132.6493
9	103.2874	36.9834	8.3002	124.0900	134.9597
10	103.3601	36.5966	8.2759	126.7540	136.9395

Table 2: Incorrectly Estimating Demand, S Model. The table displays estimates of model parameters assuming the S model. Based on generating $T = 5000$ auctions where $S = V + \varepsilon$, $V \sim N(\mu, \sigma_V^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. Here $\sigma_\varepsilon = 5$. We solve for s' as the equilibrium outcome for this game for each value of N . $\hat{s}'$ is the equilibrium outcome for the S model for each value of N .

is correctly chosen (that is a normal distribution is correctly assumed) but the S is incorrectly applied. We generated data using our Signal model for given parameters (shown in the table) and then estimated μ_V , σ_V , K and s' .

The intuition for the results in Table 2 is the opposite of that for Table 1. The S model will predict even greater selection effects than our model and so its estimated value distribution will be first order stochastically dominated by and second order stochastically dominate the distribution from the Signal model. Also, the S model should predict less entry than the Signal model and it will justify this by underestimating K .

The results in Tables 1 and 2 clearly illustrate the problems from assuming the wrong entry model. In Table 1 we overestimate μ_V and K and underestimate σ_V . The problem is exacerbated when the sample becomes “more” selected as seen in the s' column. Table 2 illustrates similar problems in estimating demand. In particular, since the model assumes that bidders have full information prior to paying their entry cost, the model underestimates the entry threshold of the true model. These biased estimates will lead to a host of problems in counterfactual analysis, which we turn to next.

Biasing Counterfactuals: Subsidies

As in any setting, working with biased estimates of demand will tend to make the economist err in counterfactual analysis. While there are many examples showing how incorrectly estimating demand may cause problems for policy recommendations, we choose to focus on one that directly stems from the selection problem which is ignored by the LS entry model: entry subsidies.

In a setting where bidders have sunk entry costs it may be attractive for a seller to subsidize the entry costs of some or all bidders (see for example Athey and Levin (2006)). In procurement auctions there may be additional policy reasons for favoring subsidies. For example, the Federal Government is mandated to buy a certain percentage of services from small businesses and it may be necessary to offer them some kind of inducement - whether in the form of an entry subsidy or

by subsidizing or favoring their bids directly - in order to meet this aim. Here, we look at a very simple example to show how estimating a model without selection could give misleading predictions about the effects of a subsidy.

To be specific, we consider a Signal model with N symmetric firms, whose values are distributed normally with $\mu = 120$ and $\sigma_V = 25$, entry costs $K = 10$, a signal with $\sigma_\varepsilon = 5$ and no reserve. With these parameters, it is expected that between 2.2 and 2.6 firms enter the auction as the number of potential entrants varies from 4 to 10. In this case there is some probability that there are no entrants or only one entrant and it is potentially attractive for the seller to encourage entry. We imagine that the seller does this by paying a subsidy to any firm that enters, and we assume that the seller can identify which firms enter (pay K) and which do not, with the aim of increasing the expected number of entrants by 1. The columns “Correct Subsidy Amount” and “Correct Δ Revenue” in Table 3 show the expected cost per bidder of the subsidy and the expected increase in revenues, respectively, when the seller uses the true model.

On the other hand, one can consider the alternative case where the seller estimates the LS model (using an infinitely large sample, so he gets the estimates reported in Table 1) and then calculates the expected cost and benefit of increasing the expected number of entrants by one using these estimates. These results appear in columns “Incorrect Subsidy Amount” and “Incorrect Δ Revenue” in Table 3, respectively. Two differences are clear. First, using the LS model results in overestimating the increase in revenues that results from additional entry because it ignores the fact that the firms who are attracted are less valuable than the inframarginal entrants. Second, it also overstates the cost of inducing additional entry. This happens because by using the LS model’s estimates, the seller believes that entry is very costly so that it is necessary to give a large subsidy to get more entry. In addition, the low estimate of σ_V reduces the amount of surplus that an additional entrant can expect to get if he wins (the expected difference in the first and second order statistics). In this example, these effects dominate the fact that in the LS model the additional entrant does not expect to be at a disadvantage to the firms that would have entered without the subsidy. Overall, the percentage biases in the revenue effects are significantly larger than the biases in the cost effects.

5 Estimation

We estimate a fully parametric model where the distributions, up to parameters, of both values and signal noise are known to the econometrician. We consider three approaches to estimation, which differ in their ability to handle heterogeneity (either in values across auctions or in the parameters across auctions and across types of bidder) and in their assumptions about equilibrium selection. We are particularly interested in whether all of the models/techniques give estimates that suggest the entry process is highly selective. In this section we only describe the estimation procedure that most will be familiar with: nested fixed point. In the appendix we describe the importance sampling and nested pseudo-likelihood approaches that we also consider.

N	$E[N]$	Correct Subsidy Amount	Incorrect Subsidy Amount	Subsidy Bias	Correct Δ Revenue	Incorrect Δ Revenue	Revenue Bias
4	2.24	9.31	13.40	-21.7%	22.28	31.75	42.69%
5	2.33	8.61	12.08	-19.4%	21.51	29.69	35.48%
6	2.40	8.17	11.26	-17.9%	20.80	28.16	30.75%
7	2.44	7.87	10.66	-16.6%	20.41	26.98	27.80%
8	2.48	7.66	10.38	-15.9%	20.26	26.54	24.40%
9	2.51	7.50	10.01	-15.1%	20.57	26.42	27.40%
10	2.53	7.37	9.77	-14.7%	20.33	26.02	24.73%

Table 3: Biased Subsidy Analysis. The table displays the effects of incorrectly assuming that the LS model is the true entry model on counterfactual analysis. The counterfactual finds the necessary reduction in K to generate an increase of one participating bidder in expectation. The table shows the correct and incorrect estimates of this subsidy. It also shows the bias in the relative changes of this subsidy and indicates that the LS model underestimates the relative change in K needed to induce one more bidder in expectation. It also shows the estimated impacts on seller revenues and the percent bias in this estimate. Again, it shows the LS model over predicts the impact on seller revenues from this subsidy. Based on generating $T = 5000$ auctions where $\mu = 120$, $\sigma_V = 25$, $\sigma_\varepsilon = 5$, $K = 10$. We solve for s' as the equilibrium outcome for this game for each value of N . The parameters used to estimate the subsidy and revenue impacts in the incorrect case appear in Table 1.

5.1 Nested Fixed Point Estimation (NFXP) using Maximum Likelihood

NFXP estimation involves solving the game for each auction at each iteration of the parameters. If there are multiple equilibria for a given set of parameters, we have to pick which is played in order to calculate the likelihood. As explained above, we assume that the parameters $\sigma_V^2, \sigma_\varepsilon^2$ and K are the same across entrant types and solve for the only type-symmetric equilibrium where the type with higher mean values has a lower s' . If other equilibria exist, we assume that they are not played.

We solve for the equilibrium entry thresholds by solving the pair of non-linear equations that define the zero profit condition for the marginal entrant of each type. For the τ type the zero profit condition is:

$$\int_0^\infty \left(v G_\tau^{-i}(v|R, N, s'_{-\tau}, s_\tau^{*'}, \theta) - \int_0^v v' g_\tau^{-i}(v'|R, N, s'_{-\tau}, s_\tau^{*'}, \theta) dv' \right) f'_\tau(v|s'_{-\tau}, s_\tau^{*'}, \theta) dv - K = 0 \quad (1)$$

where $G_\tau^{-i}(v|R, N, s'_{-\tau}, s_\tau^{*'}, \theta)$ and $g_\tau^{-i}(v'|R, N, s'_{-\tau}, s_\tau^{*'}, \theta)$ are the cdf and pdf of the highest value from other entering bidders (of both types) given their entry strategies, the number of potential entrants of each type (N), the reserve price (R) and the parameters (θ), and $f'_\tau(v|s'_{-\tau}, s_\tau^{*'}, \theta)$ is the pdf of the distribution of values for the type τ firm.

Given the equilibrium entry thresholds, we calculate the likelihood of the observed outcome of the auction. We assume that potential entrants who do not submit any bid at the auction did not enter (and so did not pay K). Because the auctions operate as open outcry auctions, rather than second-price sealed bid or button auctions, it is not entirely clear how the observed bids should be

treated. For example, the highest bid submitted by a losing bidder may be below his true value. In this case we lack a well-defined model for what determines the bid that we do observe. We therefore proceed by assuming that, when the second highest observed bid is greater than the reserve price, that this bid represents the valuation of the second highest bidder.¹⁵ Losing bidders who attend the auction are assumed to have values between this second highest bid and the reserve price. For example, the likelihood for a particular auction a when a type 1 bidder wins the auction, a type 2 bidder submits the second highest bid of b_{2a} , and $n_{\tau a}$ firms of type τ enter out of $N_{\tau a}$ potential entrants is:

$$\begin{aligned}
L_a(\theta) = & f_2(b_{2a}|\theta) * \Pr(\text{enter}_2|v_2 = b_{2a}, s_{2a}^{*'}, \theta) * \left(\int_{b_2}^{\infty} f_1(v|\theta) \Pr(\text{enter}_1|v_1 = v, s_{1a}^{*'}, \theta) dv \right) \\
& * \left(\int_R^{b_2} f_1(v|\theta) \Pr(\text{enter}_1|v_1 = v, s_{1a}^{*'}, \theta) dv \right)^{(n_{1a}-1)} * \left(\int_R^{b_2} f_2(v|\theta) \Pr(\text{enter}_2|v_2 = v, s_{2a}^{*'}, \theta) dv \right)^{(n_{2a}-1)} \\
& * \left(1 - \int_0^{\infty} f_1(v|\theta) \Pr(\text{enter}_1|v_1 = v, s_{1a}^{*'}, \theta) dv \right)^{(N_{1a}-n_{1a})} \\
& * \left(1 - \int_0^{\infty} f_2(v|\theta) \Pr(\text{enter}_2|v_2 = v, s_{2a}^{*'}, \theta) dv \right)^{(N_{2a}-n_{2a})}
\end{aligned}$$

reflecting the contributions to the likelihood of the second highest bidder, the winning bidder, the other entrants and the non-entrants, respectively. The equilibrium entry thresholds will depend on the parameters and auction characteristics, such as the reserve price which have been suppressed to reduce notation.¹⁶

5.1.1 Observed and Unobserved Auction Heterogeneity

A feature of the data is that there is considerable variation in realized prices, appraisal values (USFS predicted sale values and costs) and reserve prices across tracts. This suggests that it is important to control for observed auction heterogeneity. It may also be important to control for unobserved heterogeneity in values across auctions, since observable variables - such as the USFS's appraisal value and its estimates of logging costs and manufacturing costs, together with year and county dummies - explain less than 50% of the variation in realized revenues, and previous research (Athey, Levin, and Seira (forthcoming)) has found evidence of significant unobserved heterogeneity

¹⁵Alternative assumptions could be made. For example, we might assume that the second highest bidder has a value less than the winning bid, or that the second highest bidder's value is some explicit function of his bid and the winning bid. In practice, 96% of second highest bids are within 1% of the high bid, so that any of these alternative specifications should give similar results. We have computed some estimates using the winning bid as the second highest value and the coefficient estimates are indeed similar.

¹⁶If an entrant wins at the reserve price, then the likelihood is calculated assuming that winning bidder's value is above the reserve.

in timber auctions.¹⁷

A limitation of the NFXP algorithm is that computation time increases quickly in the number of parameters. To reduce the number of parameters we therefore attempt to control for observed heterogeneity in a first stage in a manner similar to that suggested by Haile, Hong, and Shum (2006). To do this, we first regress the log of the auction’s reserve price on observed variables including the USFS’s appraisal value for the tract and its estimates of logging and manufacturing costs. We then normalize bids and reserve prices using the predicted values from these regressions to remove the predictable component of variation and make the auctions more homogenous.¹⁸

We also use NFXP to estimate a model that allows for normally distributed *unobserved* heterogeneity in mean values across auctions. To be precise, we assume that for a given auction a $\mu_{1,a} = \mu_1 + v_a$ and $\mu_{2,a} = \mu_2 + v_a$ where $v_a \sim N(0, \sigma_v^2)$, and μ_1, μ_2 and σ_v^2 are included in the set of parameters to estimate. The parameter σ_v^2 reflects the importance of unobserved heterogeneity, which is assumed to be uncorrelated (conditional on the regressors included in the first stage) with reserve prices or the number of potential entrants. We integrate out the one-dimensional heterogeneity using a set of grid points, resulting in a simulated maximum likelihood estimator.

5.2 Identification

While we make parametric assumptions to estimate the model, here we consider informally what can identify the parameters of the model.¹⁹ With no unobserved auction heterogeneity, the distribution of values would be non-parametrically identified if entrants submit their values as bids and there is no selection. The entry process with signals can approach the no selection case when the equilibrium signal threshold for entry is very low. The equilibrium entry threshold falls when there are few potential entrants and the reserve price is low. With asymmetric bidder types, the weaker type, who will typically be more selected, will have a lower equilibrium entry threshold when there are no competitors of the stronger type.

If the distribution of values were identified from this type of variation in s' , then the distribution of signal noise and the level of entry costs would be identified from the amount of entry and changes in the distribution of observed bids (values) as s' varies due to differences in the reserve price and the number (and type) of potential entrants across auctions. For example, if signal noise is very precise then the distribution of values of entrants will be almost perfectly truncated at s' . On the other hand, if signals are very imprecise, because σ_ε is large, then the distribution of values among

¹⁷Athey, Levin, and Seira (forthcoming) estimate the distribution of unobserved heterogeneity in a first stage where they estimate a parametric bid distribution. This approach is aided by observing multiple informative bids per auction in sealed bid auctions. In open outcry auctions, it is less plausible that bids except the second-highest or winning bid should be treated as reflecting a bidder’s value.

¹⁸This will only be valid if reserve prices, which we use as the dependent variable, are set non-strategically, which is a standard assumption in the timber auction literature. One could control for observed heterogeneity in other ways, such as normalizing directly by some combination of the USFS’s predicted sale value and its measures of costs. For example, Marmer, Shneyerov, and Xu (2010) use the engineer’s cost estimate for their lawn mowing contracts. In practice, sale values and logging costs seem to have greater predictive power for both realized prices and reserve prices than other cost estimates, so we prefer to use a regression based approach which naturally allows some factors to be more important than others.

¹⁹Marmer, Shneyerov, and Xu (2010) and Einav and Esponda (2008) discuss identification in first price auctions.

entrants will be more similar to the underlying distribution of values in the population. K will be identified from the probability of entry (higher K will reduce entry) and, because of the zero profit condition, the amount of surplus the marginal entrant can expect in the auction if it enters.

As this discussion makes clear, identification of the parameters depends on having reliable measures of the number of potential entrants. In future revisions, we will look more carefully at how our estimates depend on how the set of potential entrants is defined. The discussion also assumes that firms bid their values and there is no unobserved cross-auction heterogeneity. In our data, auctions operate as open outcry auctions, so it is unreasonable to treat all of the observed bids as values. For example, losing bidders may have values above the highest bids that they announce (see Haile and Tamer (2003)). Also, unobserved heterogeneity may be important. For these reasons it is necessary to take a fully parametric approach, even though there is a lot of variation in reserve prices and the number of potential entrants.

6 Empirical Application

We now illustrate how ignoring selection into auctions matters and how to implement our methodology in an empirical setting. We focus on federal auctions of timberland in California. We first describe the data, then discuss why our model applies to the setting, then provide evidence of selection in these auctions and finally present our estimation results.

6.1 Data and Context

We focus on federal auctions of timberland in California.²⁰ In these auctions the USFS sells logging contracts to individual bidders who may or may not have manufacturing capabilities (mills and loggers, respectively). When the sale is announced, the USFS provides its own “cruise” estimate of the volume of timber on the tract. It also announces a reserve price and bidders must submit a bid of at least this amount to qualify for the auction. After the sale is announced, bidders perform their own private cruises of the tract to assess its value.²¹ These cruises can be informative about the tract’s volume, species make-up and timber quality. Finally, bidders must post a deposit of 10% of the appraised value of the tract in order to be eligible to participate in the auction.

As in our model above, we assume that bidders have independent private values. This assumption is also made in other work with similar timber auction data (see for example Baldwin, Marshall, and Richard (1997), Brannman and Froeb (2000), Haile (2001) or Athey, Levin, and Seira (forthcoming)). A bidder’s private information is primarily related to its own contracts to sell the harvest, inventories and private costs of harvesting and thus is mainly associated with its valuation only.

In our model we allow bidders to receive an imperfect signal about their value prior to paying an investment cost to fully learn their value. There are multiple reasons why it is likely that in these

²⁰We are very grateful to Susan Athey, Jonathan Levin and Enrique Seira for sharing their data with us.

²¹From our discussions with industry sources, it is very rare for firms to bid without doing their own cruise.

timber auctions bidders only have imperfect knowledge about their value prior to entry. First, each tract is unique and therefore even if a bidder has previously bid on apparently similar tracts, they must still account for heterogeneity not realized prior to further investigation. Second, the cruise estimates published by the USFS may be imprecise. We provide evidence of this below. Finally, bidders must also devote time to planning and organizing their team of harvesters and lining up potential end users for any given tract. These are likely to be only a few of the necessary investments a bidder must make prior to learning its true value for a particular stand of timber. Therefore, a model which at least allows bidders to have a noisy signal about their value, but can still permit this signal to be fairly precise, seems warranted.

We should note here that our selection model differs from Athey, Levin, and Seira (forthcoming) who use similar data but apply the LS model so that bidders have no signal about their value at the entry stage. They allow for two types of bidders (mills and loggers) and their model assumes that mills' value distribution stochastically dominates loggers', according to a hazard rate order. Their proposition 1 states that if the necessary condition for a unique type-symmetric equilibrium is met, then mills must enter with probability 1.²² In estimation the authors confirm that the necessary condition is met for each tract in their data.

From the original data we focus on the most appropriate auctions. As our methodology so far addresses ascending or second price auctions we eliminate all sealed bid auctions. We also eliminate small business set aside auctions, salvage sales and auctions with extremely low or high acreage or volume to acreage ratios as these are likely either outliers or coding errors in the data. Finally, we examine auctions between 1982 and 1989 to reduce resale concerns that might complicate the analysis (see Haile (2001) for an analysis of these auctions with resale). Resale was limited after 1981 because third party transfers (i.e. the winner transferring the right to harvest the timber) were prohibited and speculative bidding was reduced due to shortened contract lengths, larger required deposits, greater penalties for default and increased difficulty of obtaining contract extensions (Mead, Schniepp, and Watson (1983)). This is important because another model with resale like that of Haile (2001) could lead to increased bidding with increased competition, a comparative static also consistent with a selection model. We are left with 988 auctions over this period.

In addition to the USFS data, we add data on (seasonally adjusted, lagged) monthly housing starts and establishment locations. The firm location data (NETS data) was purchased from Walls and Associates who obtained the data from Dun and Bradstreet. For bidders we can identify in the NETS data, we obtain their latitude and longitude and this enables us to determine their distance from any auction. We are able to match 43.3% of firms but these firms account for 70.5% of bids and 71.0% of winning bids.

We summarize the data in Table 4. Bids are given in \$/mbf (1983 dollars). We see that bids submitted by loggers tend to be lower than those submitted by mills, consistent with the results in Athey, Levin, and Seira (forthcoming), but in our sample loggers win less often than in theirs. We define entrants to be the set of bidders we observe at the auction even if they did not submit a

²²The other alternative is that loggers enter with probability 0. This is robustly rejected in their data.

bid above the reserve price. We count the number of potential entrants as those bidders who bid within 50 km of an auction over the next month. One way of assessing the appropriateness of this definition is that less than 2% of the bidders in any auction fail to bid in another auction within 50 km of this auction over the next month. For our estimation we assume that all entrants paid the fixed cost of entry.

Fewer loggers than mills enter on average and they are also less likely to enter. Among the set of potential logger entrants, on average 34% enter, whereas on average 66% of potential mill entrants enter. Finally, 4% of tracts failed to sell because they received no bids.

Our model assumes that within type, bidders may have heterogeneous values, but not entry costs. If we believe that bidders do receive noisy signals about their value prior to entry, we can appeal to the panel nature of our data set to investigate whether there appears to be large, persistent differences in entry costs across bidders. This is because large differences in entry costs would cause those bidders with high entry costs to stay out of auctions unless they receive quite an optimistic signal about their value, in which case they will be likely to win. In the data, then, we can examine whether bidders who rarely enter are more likely to win when they do enter. We begin by focusing on bidders who are often within a reasonable distance from any auction and divide these into two groups based on how often they bid. We then will see whether the group which rarely bids is more likely to win than those who bid often. If this is true, then it is suggestive of persistence in entry costs as those bidders who rarely bid do so because of high entry costs, not low values. To proceed, we isolate a set of bidders to be those who on average are within 76.5 km of an auction, the median distance in the data, and look at the probability of winning conditional on a bidder bidding in more than 15 auctions compared to those who bid in between 5 and 15 auctions. We find that the former win 61.2% of the time they enter and the latter wins only 14.5% of the time they enter. Therefore, this suggests that there is not a great deal of persistent heterogeneity in entry costs across bidders of a certain type.

We stated earlier that government cruise reports can be inaccurate and thus bidders are incentivized to invest in learning their true value for the tract. We can support this claim in our data because for a sample of “scaled” sales we have data on the timber that was actually cut by the winner.²³ Using this information, we can evaluate the quality of the government’s estimates. On average, the government overestimates the amount of timber on the tract 60% of the time. The top panel of Figure 3 displays the distribution of these incorrect estimations in percentage terms.

It is possible that what matters more to bidders is the government’s estimate of the distribution of species type on any tract. One way to gauge their accuracy along this dimension is to compare the predicted and actual share of timber for the species that the USFS said was most populous. The bottom panel of Figure 3 displays the distribution of these incorrect estimations in percentage terms. To interpret the figure, ‘10’ means that the share of the (supposedly) most prevalent species was estimated to be 10% higher than it actually was.

Given the potential error in the government’s estimates, and the inconsistency of this error, the

²³This is the same data used in Athey and Levin (2001).

Variable	Mean	Std. Dev.	25 th -tile	50 th -tile	75 th -tile	Min	Max	N
WINNING BID (\$/mbf)	92.54	150.43	39.22	69.89	125.01	2.04	4255.73	945
BID (\$/mbf)	77.46	92.82	29.5	58.37	106.36	2.04	4255.73	3944
LOGGER	64.98	58.01	23.14	48.37	90.53	2.04	723.96	1072
MILL	82.11	102.46	32.55	62.25	113.55	5.29	4255.73	2872
LOGGER	0.27	0.44	0	0	1	0	1	3944
LOGGER WINS	0.18	0.38	0	0	0	0	1	945
FAIL	0.04	0.20	0.00	0.00	0.00	0.00	1.00	988
ENTRANTS	3.99	2.47	2	4	5	0	12	988
LOGGERS	1.09	1.32	0	1	2	0	10	988
MILLS	2.91	1.89	1	3	4	0	10	988
POTENTIAL ENTRANTS	9.98	6.45	5	9	14	1	38	988
LOGGER	5.26	4.71	2	4	8	0	27	988
MILL	4.73	2.78	3	4	6	0	14	988
PREVIOUS MILLS (6 mos)	4.57	2.85	2	4	7	0	14	988
SPECIES HERFINDAHL	0.55	0.23	0.35	0.51	0.72	0.20	1.00	988
DENSITY (acres/mbf)	0.21	0.21	0.07	0.16	0.28	0.02	1.81	988
VOLUME (hundred mbf)	75.17	45.41	41.05	68.6	103.1	5	275.4	988
HOUSING STARTS	1610.74	267.8	1580.5	1628	1782	843	2260	988
RESERVE (\$/mbf)	37.95	32.36	16.38	27.3	47.89	2.04	221.87	988
SELL VALUE (\$/mbf)	291.05	64.18	259.51	292.29	326.03	0	518.95	976
ROAD CONST (\$/mbf)	12.23	14.33	1.06	7.49	17.76	0	91.55	976
LOG COSTS (\$/mbf)	116.53	33.13	98.4	113.12	133.78	0	252.46	976
MFCT COSTS (\$/mbf)	134.46	24.09	127.07	136.22	146.15	0	227.55	976
MISSING APPRAISAL	0.01	0.11	0	0	0	0	1	988
DISTANCE (KM)	86.91	122.58	37.7	60.2	86.48	0.86	1017.84	2702

Table 4: Summary statistics for California ascending auctions from 1982-1989. We exclude SBA set asides, salvage sales, auctions with very high or low volume to acreage ratios and failed sales. For calculating when a logger wins, we focus only on auctions where the tract sold. We count the number of potential entrants as those bidders who bid within 50km of an auction over the next month. Statistics about the potential mill bidders over the previous 182 days that bid in the same forest district is also given by PREVIOUS MILLS. We note that 26% of all bids are losing bids at the reserve. SPECIES HERFINDAHL is the Herfindahl index for wood species concentration on the tract. SELL VALUE, ROAD CONST, LOG COSTS and MFCT COSTS are USFS appraisals of the value of the tract and the road building, logging and manufacturing costs of the tract, respectively. HOUSING STARTS is the seasonally adjusted, lagged monthly housing starts in a tract's county. DISTANCE is the straight line distance between a bidder's establishment location and a tract's centroid.





Figure 3: Evaluating the quality of the government's predictions.

government’s estimates are useful but not perfect in describing these tracts. Therefore, bidders likely find it valuable to undertake their own investment to more precisely learn about their value. Moreover, our discussions with industry sources confirm that it would be extremely unlikely for bidders to participate without previously undertaking their own cruise.

6.2 Evidence of Selection

In this section we provide evidence that actual bidders in an auction are not a random sample of potential entrants from the data. We do this by examining the impact of the number of potential entrants on submitted bids.

One test for selection is whether the average valuations of bidders rise as the number of potential entrants increase. If there is no selection, then bidders are a random sample from the population regardless of the number of potential entry. This is the essential idea behind the test in Marmer, Shneyerov, and Xu (2010). If we believe that bidders’ strategies in the auctions examined here are the same as in a classic English Button Auction with independent private values, then we can consider bids submitted to be bidders’ valuations. Thus, we can examine whether valuations increase in potential entry by looking at the submitted bids as potential entry increases. These results appear in Table 5. Examining the first four columns, we find that a 10% increase in potential entry leads to a 1.6% increase in the submitted bid (0.9% when we control for distance), thus providing evidence of selection. A confounding factor is that there may be unobserved auction heterogeneity that is driving both increased potential entry and submitted bids. That is, there may be factors observable to the bidders but not to the econometrician that affect bidder behavior. To control for this we follow a strategy similar to Haile (2001) and instrument for the number of potential bidders by the number of mills who bid in the same forest during the preceding six months. If we believe mill location and activity are determined well before a particular auction’s unobservable (to the researcher) quality is realized, then this is a valid instrument. The estimates employing this instrument appear in the last four columns. We continue to find a positive impact of potential entry on submitted bids. In fact, the impact increases by almost three-fold.^{24,25} While there may be some concern about this instrumental variable strategy, in our structural estimation we do not rely on this method to control for unobserved heterogeneity. Here we are simply interested in providing suggestive evidence that there is selection at the entry stage for these timber auctions.

There is some concern about interpreting bids as values in these auctions since they are not exactly English Button Auctions, but rather open outcry auctions (see for example Haile and Tamer (2003)). If we assume that any bid submitted is a value, it is safest to assume that the highest

²⁴If we include the reserve price, following the arguments made in Roberts (2009), to control for unobserved heterogeneity, we also find evidence of selection. We are less enthusiastic about this route since it seems that the reserve pricing function used by the USFS may not always lead to reserve prices being monotonic in unobservable quality (Haile and Tamer (2003)).

²⁵One might not expect the magnitude to increase. Two potential explanations for this effect are that we are actually instrumenting for the number of mills, and an additional potential mill entrant is more valuable than an additional potential logger entrant, and measurement error. In future revisions we aim to more clearly parse these potential explanations’ effects.

	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log POT. ENTRANTS	0.157*** (0.024)	0.155*** (0.024)	0.083*** (0.022)	0.092*** (0.026)	0.361*** (0.077)	0.381*** (0.078)	0.221*** (0.076)	0.281*** (0.09)
LOGGER	-.159*** (0.029)	-.145*** (0.029)	-.121*** (0.027)	-.100*** (0.032)	-.166*** (0.029)	-.154*** (0.03)	-.125*** (0.027)	-.101*** (0.032)
SCALE		0.246*** (0.061)	0.178*** (0.056)	0.082 (0.067)		0.231*** (0.061)	0.167*** (0.056)	0.064 (0.067)
SPECIES HERFINDAHL		-.663*** (0.063)	-.367*** (0.06)	-.359*** (0.071)		-.675*** (0.063)	-.381*** (0.06)	-.379*** (0.071)
DENSITY (acres/mbf)		-.052 (0.071)	0.107 (0.066)	0.039 (0.079)		0.0001 (0.074)	0.135** (0.067)	0.096 (0.083)
VOLUME (mbf)		-.0007** (0.0003)	-.001*** (0.0003)	-.0007* (0.0003)		-.0007** (0.0003)	-.001*** (0.0003)	-.0006* (0.0003)
HOUSING STARTS		0.0002** (0.00009)	0.0002** (0.00008)	0.0003*** (0.0001)		0.0001 (0.00009)	0.0002* (0.00009)	0.0003*** (0.0001)
log SELL VALUE (\$/mbf)			1.418*** (0.066)	1.462*** (0.08)			1.367*** (0.071)	1.393*** (0.086)
log ROAD CONST (\$/mbf)			0.057*** (0.011)	0.04*** (0.013)			0.059*** (0.011)	0.043*** (0.013)
log LOG COSTS (\$/mbf)			-1.536*** (0.063)	-1.617*** (0.078)			-1.488*** (0.068)	-1.554*** (0.083)
log MFCT COSTS (\$/mbf)			-.061 (0.065)	-.111 (0.084)			-.047 (0.065)	-.089 (0.084)
MISSING APPRAISAL			0.411** (0.186)	0.034 (0.215)			0.395** (0.186)	0.013 (0.215)
log DISTANCE (KM)				-.001 (0.019)				-.005 (0.019)
YEAR DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
QUARTER DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
COUNTY DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
FIRST STAGE					0.079*** (0.004)	0.077*** (0.004)	0.075*** (0.004)	0.075*** (0.005)
IV COEFFICIENT								
N	3944	3944	3944	2702	3944	3944	3944	2702
R ²	0.227	0.256	0.379	0.382				

Table 5: Impact of potential entrants on all bids. Dependent variable is log of bid per volume. POT. ENTRANTS are the bidders who bid in this auction or in those within 50 km over the next month. Instrument for potential entrants is the number of unique mill bidders in a forest during the previous 6 months. Observations drop when we include distance because we are missing that data for some bidders.

bid is the valuation of the second highest valuation bidder. Therefore, to further push on whether we can find evidence of selection, we repeat the same test using only the winning bid. The results appear in Table 6 and corroborate the evidence just presented using all bids.

Another way to investigate the amount of selection in the data is to analyze the impact of potential entry on revenues. We know from above that models without selection will predict that revenues fall as potential entry increases. Given that so few auctions fail to sell, the results are virtually identical to those in Table 6, thus lending support to a model with selection.

Finally, we can test for evidence of selection by applying a Heckman selection model (Heckman (1976)). Given that potential competition affects a bidder's decision to enter an auction, but not his value conditional on entry, we can specify entry as a flexible function of covariates, including potential competition, at the first stage, and exclude potential competition at the second stage bid regression. A related method is used in Ellickson and Misra (forthcoming) in analyzing supermarket pricing techniques when *ex post* revenue data is available. The results appear in Table 7. In the first stage entry probit, we include potential other mill and logger entrants and incorporate them through a flexible polynomial. It is clear that we are finding positive selection in that entrants are likely to bid more aggressively. In addition, we see that the difference in logger and mill bids is masked by selection because the difference grows when we control for selection. The results are robust across a variety of specifications and restricting the sample to only winning bids.²⁶

After establishing the evidence of selection in the data, we now turn to estimating the full entry and bidding model in order to perform counterfactual analysis.

6.3 Estimates

We now present estimates of the structural signal model. To estimate the model we use 888 out of the 988 auctions used above. This selection is done primarily to drop outliers. For example, a few auctions have very low USFS estimated sale values relative to the other auctions, and the highest winning bid is over \$4,255/mbf whereas the 99th percentile is only \$320/mbf. Specifically to be in the estimation sample an auction must have:

1. a winning bid between \$5/mbf and \$350/mbf or result in no sale.
2. a USFS estimate of the expected sale value of the timber between \$184/mbf and \$428/mbf.
3. a USFS estimate of logging and manufacturing costs.
4. no more than 20 potential entrants (using the definition discussed above).
5. a location that can be identified so that the number of potential entrants can be calculated.

As can be seen from these definitions, considerable amounts of heterogeneity across auctions remains. All of the specifications assume that values are distributed log normal.

²⁶ Although not shown in the table, the results are also robust to changing the regressions from logs to levels.

	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log POT ENTRANTS	0.279*** (0.037)	0.273*** (0.037)	0.19*** (0.032)	0.174*** (0.039)	1.119*** (0.144)	1.130*** (0.147)	0.924*** (0.136)	1.086*** (0.203)
LOGGER	-.222*** (0.07)	-.192*** (0.071)	-.074 (0.062)	-.026 (0.076)	-.280*** (0.086)	-.285*** (0.089)	-.155** (0.077)	-.111 (0.103)
SCALE		0.416*** (0.119)	0.32*** (0.102)	0.169 (0.126)		0.327** (0.148)	0.246* (0.126)	0.071 (0.168)
SPECIES HERFINDAHL		-.618*** (0.113)	-.332*** (0.102)	-.368*** (0.124)		-.598*** (0.14)	-.360*** (0.125)	-.340** (0.165)
DENSITY (acres/mbf)		-.066 (0.12)	0.065 (0.103)	-.037 (0.127)		0.117 (0.152)	0.199 (0.129)	0.294 (0.182)
VOLUME (mbf)		-.013** (0.005)	-.015*** (0.005)	-.011* (0.006)		-.019*** (0.007)	-.020*** (0.006)	-.012 (0.007)
HOUSING STARTS		0.0002 (0.0002)	0.0002 (0.0001)	0.0002 (0.0002)		0.00007 (0.0002)	0.00003 (0.0002)	0.00009 (0.0002)
log SELL VALUE (\$/mbf)			1.276*** (0.102)	1.135*** (0.119)			0.952*** (0.139)	0.843*** (0.17)
log ROAD CONST (\$/mbf)			0.047*** (0.018)	0.03 (0.022)			0.065*** (0.022)	0.05* (0.029)
log LOG COSTS (\$/mbf)			-1.766*** (0.1)	-1.552*** (0.122)			-1.437*** (0.136)	-1.242*** (0.174)
log MFCT COSTS (\$/mbf)			0.315*** (0.096)	0.229* (0.125)			0.376*** (0.119)	0.282* (0.165)
MISSING APPRAISAL			0.131 (0.317)	-.113 (0.399)			-.052 (0.391)	-.187 (0.527)
log DISTANCE (KM)				0.076** (0.032)				0.027 (0.044)
YEAR DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
QUARTER DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
COUNTY DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
FIRST STAGE IV COEFFICIENT					0.094*** (0.009)	0.092*** (0.009)	0.087*** (0.009)	0.075*** (0.012)
N	945	945	945	671	945	945	945	671
R ²	0.36	0.396	0.563	0.551				

Table 6: Impact of potential entrants on winning bids. Dependent variable is log of the winning bid per volume. POT. ENTRANTS are the bidders who bid in this auction or in those within 50 km over the next month. Instrument for potential entrants is the number of unique mill bidders in a forest during the previous 6 months. Observations drop when we include distance because we are missing that data for some bidders.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CONSTANT	2.941*** (0.421)	2.690*** (0.421)	4.030*** (0.785)	3.071*** (0.768)	3.335*** (0.424)	3.333*** (0.423)	3.043*** (0.72)	3.037*** (0.713)
LOGGER	-.154*** (0.029)	-.388*** (0.047)	-.202*** (0.072)	-.638*** (0.096)	-.118*** (0.027)	-.211*** (0.041)	-.054 (0.063)	-.285*** (0.081)
SCALE					0.184*** (0.056)	0.187*** (0.056)	0.34*** (0.104)	0.339*** (0.103)
SPECIES HERFINDAHL					-.359*** (0.06)	-.351*** (0.06)	-.324*** (0.104)	-.292*** (0.103)
DENSITY (acres/mbf)					0.09 (0.066)	0.081 (0.066)	0.031 (0.105)	0.0008 (0.104)
VOLUME (mbf)					-.001*** (0.0003)	-.001*** (0.0003)	-.001*** (0.0005)	-.002*** (0.0005)
HOUSING STARTS					0.0002** (0.00008)	0.0002** (0.00008)	0.0002 (0.0001)	0.0001 (0.0001)
log SELL VALUE (\$/mbf)					1.448*** (0.066)	1.466*** (0.066)	1.360*** (0.104)	1.397*** (0.103)
log ROAD CONST (\$/mbf)					0.055*** (0.011)	0.051*** (0.011)	0.042** (0.018)	0.035* (0.018)
log LOG COSTS (\$/mbf)					-1.565*** (0.063)	-1.608*** (0.065)	-1.851*** (0.101)	-1.946*** (0.102)
log MFCT COSTS (\$/mbf)					-.070 (0.065)	-.049 (0.065)	0.299*** (0.098)	0.349*** (0.098)
MISSING APPRAISAL					0.421** (0.186)	0.437** (0.186)	0.178 (0.323)	0.221 (0.32)
$\hat{\lambda}$		0.342*** (0.054)		0.658*** (0.098)		0.136*** (0.046)		0.341*** (0.077)
YEAR DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
QUARTER DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
COUNTY DUMMIES	YES	YES	YES	YES	YES	YES	YES	YES
ONLY WIN BID	NO	NO	YES	YES	NO	NO	YES	YES
N	3944	3944	945	945	3944	3944	945	945
R ²	0.219	0.227	0.32	0.352	0.377	0.378	0.546	0.555

Table 7: Heckman selection evidence. Dependent variable is log of the bid per volume. The table displays the second stage results of the two step selection model. The first stage probit is of entry where the exogenous shifters are potential other mill and logger entrants, incorporated as a flexible polynomial. $\hat{\lambda}$ is the estimated inverse Mills ratio from the first stage.

In all specifications we assume that values are distributed lognormal and that the signal noise is multiplicative, i.e., $S_\tau = V_\tau A_\tau$, $A_\tau = e^{\varepsilon_\tau}$, where $V_\tau \sim \log N(\mu_\tau, \sigma_{V_\tau}^2)$, $\varepsilon_\tau \sim N(0, \sigma_{\varepsilon_\tau}^2)$, where $\tau \in \{\text{Mill, Logger}\}$. We have also estimated several models where values are normally distributed. These estimates provide even stronger evidence of selection, but have the unattractive feature that they imply that a significant proportion of potential entrants have negative values. In future revisions, we will estimate models where values can take on more flexible distributions such as Weibull, and we will also estimate models where the distribution of values is truncated at the reserve price.²⁷

While we present several sets of parameter estimates (those for the importance sampling and nested pseudo likelihood approaches are in appendices), our main findings about value distributions and the degree of selection are summarized in Table 8, for an auction with average characteristics including 4 mill and 5 logger potential entrants and a reserve price of \$37/mbf. Each row of the table reflects a different empirical specification. The first three columns report the mean value for the average mill, the average mill that enters and the marginal mill (i.e., the expected value for a mill receiving s'). The difference between these columns reflects the degree of selection among mills. The next set of columns report mean values for loggers. The difference between the average values of mills and loggers indicates the degree of asymmetry between bidder types. A feature for all of the specifications is that the marginal mill has a lower expected value than the average mill. This simply reflects the fact that most mills enter. On the other hand, the marginal logger has a higher value than the average logger, because most loggers choose not to enter. For the importance sampling estimates we also display a picture of the value distributions for all, entrant and marginal bidders, for both mills and loggers. The figures are similar for all estimates and so for brevity sake we do not include them here.

6.3.1 Nested Fixed Point Estimates

Table 9 presents the results from three NFXP specifications. In all cases we control for observable factors that affect values using a pre-estimation regression as explained above. The specifications differ in their allowance for unobserved heterogeneity and the definition of the likelihood. The estimates in the first and third columns assume that there is no unobserved heterogeneity, while in the second column we allow for unobserved heterogeneity in values that is normally distributed and common to mills and loggers. The first and second specifications assume that non-bidders did not pay K ; the third specification does not. We describe how this alters the likelihood below.

All sets of estimates imply that the asymmetry between the average values of mills and loggers is much larger than the difference in bids (\$20/mbf or 21%), which would be the basis of an estimate that ignored selection. The estimate of σ_η in the second column indicates that, controlling for observables, there is evidence of significant heterogeneity in mean values across auctions. In

²⁷Atthey, Levin, and Seira (forthcoming) assume that the distribution of values is approximately truncated at the reserve price in timber auctions. This requires entry costs to be relatively high to explain why firms do not enter. In contrast, our current estimates imply that in many auctions potential entrants have values below the reserve price so that many potential entrants do not want to enter if signals are relatively precise even when K is small.

Method	Table	Mill			Logger		
		All	Entrant	Marginal	All	Entrant	Marginal
NFXP	9	69.91	95.98	30.67	29.63	66.87	32.71
NFXP (Unobs. Hetero.)	9	79.83	126.47	36.79	28.46	82.55	37.33
NFXP (LL#2)	9	72.83	94.88	24.78	26.56	58.70	26.24
NPL (Same)	14	62.00	88.54	30.61	23.00	59.17	31.27
NPL (Different)	14	61.51	95.86	34.74	30.64	63.65	37.16
Importance Sampling	11	71.81	110.07	36.50	24.68	74.93	38.87
Importance Sampling (More Params and LL#2)	12	71.64	85.75	22.47	18.70	46.29	24.59

Table 8: Summary of parameter estimates. The table displays estimates of the mean values for all, entrant and marginal mills and loggers in an auction with a reserve of \$37/mbf, 4 potential mill entrants and 5 potential logger entrants. The importance sampling estimates display the median value of the estimates. LL#2 indicates that we are allowing non-bidders to have entered and learned that their value was less than the reserve.

	(1)	(2)	(3)
μ_{mill}	3.773 (0.018)	3.783 (0.024)	3.759 (0.019)
μ_{logger}	2.913 (0.017)	2.749 (0.024)	2.748 (0.019)
σ_V	0.975 (0.001)	1.094 (0.001)	1.030 (0.002)
σ_ε	0.3098 (0.0088)	0.0328 (0.010)	0.2954 (0.025)
K	0.0842 (0.0100)	0.0110 (0.002)	0.0155 (0.004)
σ_η	-	0.625 (0.002)	-
N	888 auctions	888 auctions	888 auctions
Log-likelihood	-11,085.6	-10,738.2	-10,694.6
Assume only entrants pay K ?	YES	YES	NO

Table 9: NFXP parameter estimates. Nested fixed point estimates of parameters for lognormal parameterization. The specifications differ by their allowance for unobserved heterogeneity and the definition of the likelihood. The second specification allows for there to be unobserved heterogeneity in values which is normally distributed and common to mills and loggers. The first and second specification make the assumption that non-bidders did not pay K . Standard errors are in parentheses.

comparing estimates of σ_V in the first two columns, we note that currently allowing for heterogeneity across auctions does not reduce the estimated dispersion of values within a given auction, but that it instead increases the estimated degree of selection (precision of ε), which is a second factor that, all else equal, can reduce the dispersion of observed bids.²⁸ We note that the the average value for mills in this model is higher than that in other models. We believe that this reflects the fact that our normality assumption for the distribution of unobserved heterogeneity is not quite appropriate for this set of auctions because it implies that in a small number of auctions values should be exceptionally high, outside the range that we observe in the data. In future revisions we will therefore try some alternative assumptions.

The final column of Table 9 contains estimates where we change our definition of the likelihood. The previous estimates assume that firms that do not submit qualifying bids do not enter (i.e., they do not pay K). However, this rules out the possibility that they enter, but then do not attend the auction because they discover that their value is less than the reserve price. The change to the likelihood is to make the probability that a type τ potential entrant does not attend the auction equal to $(1 - \int_R^\infty f_\tau(v|\theta) \Pr(\text{enter}_\tau|v_\tau = v, s_{\tau a}^*, \theta) dv)$. The change maintains the estimated difference between logger and mills but, as expected, lowers the estimate of K .

The NFXP estimates assume that, if there are multiple equilibria, the firms will play the equilibrium where mills have the lower s' . We now check whether the parameters can support multiple equilibria. Figure 4 shows the type-symmetric “equilibrium best response functions” for mills and

²⁸A closer look at the data indicates that it may be necessary to allow the degree of unobserved heterogeneity to vary with the USFS’s appraisal value, as auctions with lower appraisal values appear to be more heterogeneous.

loggers for the representative average auction considered in Table 8, based on the parameters in column (1). By an “equilibrium best response function” for mills we mean the strategy that would represent a symmetric equilibrium among the set of 4 mills if the 5 loggers all used a particular s' . As usual, type-symmetric Bayesian Nash Equilibria are given by combinations where the best response functions intersect. For these parameters and this auction there is only a single equilibrium due to the best response functions’ flatness. This reflects the large asymmetry between the types and the relatively small value of K . That is, firms want to enter if they are likely to have values above the reserve even if their probability of winning is relatively small. While the estimates are conditional on a particular equilibrium selection assumption, this provides some evidence that concerns about multiple equilibria may not be too important empirically. This is also true for other auctions and estimates that we have analyzed.




Figure 4: Best response functions for mills (solid) and loggers (dash-dot) based on nested fixed point parameter estimates for specification (1) in Table 9.

To give a better sense of what the distribution of values looks like, we present some examples in Figures 5-6. We plot the distributions for the results in Table 12. The value distributions for both types appear in Figure 5. The distributions corresponding to our other estimation methods are very similar.

With our parameter estimates, if we assume a set of potential bidders and reserve price, we can compare the entrants’ and marginals’ (those who observed s') value distributions. For the case where the reserve and the number of potential mill and logger entrants are set to their respective means of \$37/mbf, four and five, this comparison for mills and loggers appears in Figure 6. As




Figure 5: Comparing the value distributions for Mills (dashed) and Loggers (solid). Based on estimates in Table 12.

predicted by the Signal model, there is a stark difference in the marginal and infra-marginal bidders for each type.

7 The Benefits of Increased Competition with Entry

Bulow and Klemperer (1996) establish a landmark result on the benefits of competition. They show that sellers prefer auctions with no reserve to those in which they can set an optimal reserve but have one less bidder if (1) bidders' marginal revenue curves are downward sloping, (2) every bidder is willing to make an opening offer of at least the seller's value in an English auction, (3) bidders are symmetric and (4) participation is fixed (Bulow and Klemperer (1996)). To understand this result²⁹, observe that the seller could do as well guaranteeing sale with $N + 1$ buyers as he could running an optimal auction with N buyers: he could just run the optimal auction on an initial set of N bidders and if the item didn't sell, simply give it to the $N + 1$ th bidder. However, we know from Myerson (1981) that the English auction with independent values is the optimal mechanism among those that guarantee sale.³⁰ That is, it is strictly better than the mechanism where the seller runs the optimal auction on the first N bidders and then assigns the good to the $N + 1$ th bidder in the event of no sale.

²⁹This discussion draws on Kirkegaard (2006).

³⁰Strictly speaking, this is true as long as the buyer with the highest marginal revenue wins and any buyer with the lowest possible signal earns zero payoff from participating.




Figure 6: Comparing the value distributions for entrant and “marginal” bidders by type. Based on estimates in Table 12.

The result is not robust to including costly bidder participation where we now consider increasing the number of potential entrants from N to $N + 1$. This can trivially be seen in the LS model. Here the optimal reserve price is independent of the number of potential bidders and equal to 0. We also know that if bidders enter with probability less than one, the seller’s expected revenues decline in N . Conversely, the seller may prefer to increase the number of potential bidders by one and forgo the opportunity to set a reserve when there is selective entry. For example, in the S model, if values are distributed $\log N(5, .4)$, $K = 20$ and $N = 4$, with no reserve the seller’s expected revenues are 128.9. If instead $N = 3$ but the seller can set the optimal reserve price of 73.9, the seller’s expected revenues are 118.17. The relative efficiency of the LS and S models, from the seller’s perspective, makes the comparison of the benefits of increased competition across the models intuitive. When the number of potential entrants increases, each previous potential entrant is less likely to participate. This counteracts the benefits naturally stemming from a larger pool of potential buyers. In a model with no selection, there is a chance that “valuable” bidders to the seller cease participating, thus making the “cost” of increased potential entry nontrivial. In a model with selection, on the other hand, the bidders most likely to stop participating are the least valuable among previous entrants. This in turn minimizes the “cost” of increasing potential entry.

Therefore, when costly entry is introduced, it is an empirical question whether the seller would prefer the ability to set a reserve price or give up this power to gain an extra potential bidder. It is here that our flexible model of entry can be used to evaluate this question without imposing restrictions that will determine the answer *a priori*. Table 10 compares the predicted revenues from auctions with an optimal reserve to those without a reserve but one more bidder. The top (bottom) panel pertains to increasing the number of potential mill (logger) entrants. The first row gives the expected revenues if the seller can set the optimal reserve and faces five potential loggers and N_M potential mills, with N_M varying across columns. The optimal reserve appears at the bottom of this

<i>Vary Mills:</i>					
$N_L = 5$ and $N_M =$	1	2	3	4	5
Expected Revenues with Optimal Reserve	34.688	50.199	64.517	77.341	89.020
Expected Revenues with No Reserve but:					
$N_M + 1$ Potential Mills	48.948	63.683	76.874	88.708	99.319
Optimal Reserve	35.75	44.75	46.75	49.75	53.75
$Pr[\text{No Sale}]$ with Optimal Reserve	0.218	0.172	0.101	0.070	0.053
<i>Vary Loggers:</i>					
$N_M = 5$ and $N_L =$	1	2	3	4	5
Expected Revenues with Optimal Reserve	73.076	74.166	75.245	76.267	77.341
Expected Revenues with No Reserve but:					
$N_L + 1$ Potential Loggers	73.208	74.474	75.687	76.874	77.940
Optimal Reserve	59.50	57.50	50.50	50.00	49.75
$Pr[\text{No Sale}]$ with Optimal Reserve	0.147	0.128	0.084	0.076	0.070

Table 10: The Value of Competition. The table simulates the potential revenues for a set of bidders when the seller sets the optimal reserve to the case where he cannot set the optimal reserve but has one more bidder of a certain type. Results are based on parameter estimates in Table 12 and 100,000 simulated auctions.

panel along with the corresponding probability that no sale occurs. In the table's second row we compute the expected revenues if the seller faces one additional potential mill bidder but is unable to set a reserve (so the number of mills in each column is $N_M + 1$). The revenues rise for each of the considered cases (as they do for all cases we have tried). The increase in revenues comes despite the fact that with the additional mill all previous potential entrants increase their s' . The benefit the additional mill brings declines in the number of potential entrants, but still remains high as even with many potential entrants. For example, an additional potential mill entrant increases expected revenues by 11.5% relative to when the seller could set an optimal reserve but only had five potential mill and five potential logger entrants.

The second panel in Table 10 repeats the exercise holding the number of potential mill entrants fixed at four and increasing the number of potential logger entrants. While the gains from increased potential logger competition are not as great as that of mills, the seller always prefers an extra logger potential bidder to the ability to set an optimal reserve. Interestingly, the optimal reserve price falls in the number of potential logger entrants.

8 Conclusion

In most of the empirical auction literature, the econometrician assumes that the bids observed in the data are submitted by a random sample of bidders. In this paper we relax this assumption, developing and estimating an independent private value auction model where potential bidders receive an imperfect signal about their value prior to making a costly entry decision. The estimation allows for asymmetries in bidders and unobserved, to the econometrician, object heterogeneity.

We illustrate how incorrectly assuming that the current leading models of entry are the data

generating processes will (i) cause biased estimates of model primitives and (ii) generate bias in important counterfactual analyses. We also highlight that a feature of models without selection, namely that seller revenues decrease in the number of potential bidders, can be used to test for the presence of selection. This is because seller revenues can increase with potential entry in the more general Signal model.

We apply our model to timber auctions to demonstrate how ignoring selection in bidder participation can affect an important and often studied empirical setting. Across a variety of specifications our structural estimates suggest that potential bidders have quite precise estimates of their value prior to paying an entry cost. This precision generates a high degree of selection in the data.

We take advantage of the flexible entry model to reevaluate the well known theory result regarding the benefits of competition. Specifically, Bulow and Klemperer (1996) show that with fixed participation a seller prefers an auction with no reserve to one with an optimal reserve but one less bidder. Extending the result to statements about the effects of increasing potential competition is not straightforward. In particular, depending on the amount of selection in the entry process, the seller may not prefer to increase the number of potential entrants if he must lose the ability to set an optimal reserve price. Our flexible model of entry, therefore, can be usefully employed to understand the benefits of competition when entry is costly. Using our structural estimates we find if its goal is to maximize auction revenue, then the USFS should willingly forgo its ability to set an optimal reserve price in favor of attracting one more potential entrant.

A Importance Sampling

The models estimated using NFXP assume that there is no heterogeneity in the parameters σ_V^2 , K and σ_ε^2 across auctions. This may be incorrect as it is at least plausible that features of the tract, such as the density of wood and whether the sale is a scaled sale, would affect these parameters.³¹ To allow for heterogeneity in multiple parameters we follow the importance sampling approach suggested by Akerberg (2009). The benefit of this approach is that we do not need to resolve the entry game for each auction at each value of the parameters. Instead, we assume that all of the parameters of the model have some parametric distribution across auctions, which can depend on observables. We will assume that the parameters are jointly normally distributed with truncation where appropriate. We solve each auction and calculate the likelihood of the observed decisions for many different simulation draws of the parameters (done using many processors) and then estimate the distribution of the parameters by reweighting the likelihoods for each simulated game.³² When solving the games we assume that the equilibrium played involves the type with the higher mean having a lower s' , and to make sure that there is exactly one equilibrium of this kind we assume that the parameters σ_V^2 , K and σ_ε^2 are the same across potential entrants within an auction even if they are different across auctions.

Specifically we assume that

$$\begin{aligned} V_{mill} &\sim \log N(\mu_{mill,a}, \sigma_{V,a}^2) \text{ where } \mu_{mill,a} \sim N(X_a\beta_1, \omega_{\mu,mill}^2) \\ V_{logger} &\sim \log N(\mu_{mill,a} + \mu_{difference,a}, \sigma_{V,a}^2) \text{ where } \mu_{difference,a} \sim \text{Truncated } N(X_a\beta_2, \omega_{\mu,difference}^2) \\ \sigma_{V,a}^2 &\sim \text{Truncated } N(X_a\beta_3, \omega_{\sigma_V^2}^2) \\ K_a &\sim \text{Truncated } N(X_a\beta_4, \omega_K^2) \\ \sigma_{\varepsilon,a}^2 &\sim \text{Truncated } N(X_a\beta_5, \omega_{\sigma_\varepsilon^2}^2) \end{aligned}$$

where the X s are observable variables. The parameters to estimate are the parameters which characterize the distributions of the structural parameters i.e., the β s, $\omega_{\mu,mill}^2$, $\omega_{\mu,difference}^2$, $\omega_{\sigma_V^2}^2$, ω_K^2 and $\omega_{\sigma_\varepsilon^2}^2$. We label the collection of these parameters Ω . The truncation points are chosen so that loggers always have weakly lower mean values than mills and both K and the variance parameters are positive.

For given parameters Ω the likelihood function for the observed outcome in auction a is

$$\int L_a(\theta) \phi(\theta|X_a, \Omega) d\theta$$

³¹In a scaled sale a seller bids a price for a given volume of wood and the total price paid is determined *ex post* by the total amount of wood on the tract. A bidder's value may therefore depend less on the total volume of wood on the tract. On the other hand, in a non-scaled sale a bidder submits a price for all of the wood on the tract and so it is much more important to estimate the amount of wood on the tract accurately.

³²Hartmann (2006) and Hartmann and Nair (forthcoming) provide applications of these methods to consumer dynamic discrete problems. Bajari, Hong, and Ryan (forthcoming) use a related method to analyze entry into first price procurement auctions assuming the LS model. However, their method assumes a complete information entry game where firms get draws on entry costs that depend on who else enters the auction. We do not make this assumption.

where $L_a(\theta)$ is the likelihood of the outcome for a given auction for a given set of structural parameters defined in the discussion of the NFXP algorithm. Estimation becomes costly if whenever one alters Ω the set of θ s to be evaluated is also altered. To motivate the use of importance sampling this can be rewritten as

$$\int L_a(\theta) \frac{\phi(\theta|X_a, \Omega)}{g(\theta|X_a)} g(\theta|X_a) d\theta$$

where $g(\theta|X_a)$ is an importance sampling density where the support of θ does not depend on Ω , which is true in our case because the truncation points are not functions of the parameters. The likelihood can be simulated using

$$\frac{1}{S} \sum_s L_a(\theta_s) \frac{\phi(\theta_s|X_a, \Omega)}{g(\theta_s|X_a)}$$

where θ_s are vectors of parameter draws from $g(\theta_s|X_a)$. For g we use uniform distributions over a very wide area of the parameter space with 80,000 draws for each auction. During estimation the weights $\frac{\phi(\theta_s|X_a, \Omega)}{g(\theta_s|X_a)}$ change as Ω varies but $L_a(\theta_s)$ does not have to be recalculated. This allows us to control for a large number of observable variables directly, rather than using an ad-hoc first stage. Standard errors are calculated using a non-parametric bootstrap where we resample games and their associated draws with replacement.

A.1 Importance Sampling Results

The importance sampling estimates shown in Table 11 make the same equilibrium assumptions as the NFXP estimates, but allow us to control for observed auction heterogeneity as part of estimation and to allow for unobserved heterogeneity in the parameters across auctions. The estimates in Table 11 allow the USFS estimate of sale value and its estimate of logging costs to affect mill and logger values (where $\mu_{\text{logger}} = \mu_{\text{mill}} + \mu_{\text{difference}}$). These variables are consistently the most significant in regressions of reserve prices or winning bids on observables. The right hand column shows the median value of the parameters across the 888 auctions in the sample, taking into account observable auction characteristics and the fact that some of the parameter distributions are truncated. The coefficients show that tracts with greater sale values and lower costs are more valuable as one would expect. Interestingly the difference in values between mills and loggers appears independent of these variables, although it does appear that there is both unobserved heterogeneity in values across auctions (the standard deviation of μ_{mill}) and heterogeneity in the difference between mill and logger mean values across auctions (the standard deviation of $\mu_{\text{difference}}$), which was not allowed for in any of the NFXP estimates.

The estimates for K and σ_ε in Table 11 indicate that entry costs are low and signals are quite precise, leading to a lot of selection.

The specification in Table 12 adds additional observable heterogeneity to, and uses the different definition of the likelihood, than the specification in Table 11. Here we allow for the possibility that potential entrants that do not attend the auction paid K , but learned that their values were

	Constant	Log(Sale Value)	Log(Logging Costs)	Std. Deviation	Median Value of Parameter
μ_{mill}	-8.4653 (0.8433)	3.1630 (0.2042)	-1.1968 (0.1525)	0.3116 (0.0275)	3.7688
$\mu_{difference}$	-1.0508 (0.5725)	-0.0008 (0.0021)	-0.0015 (0.0026)	0.4046 (0.4267)	-1.0701
σ_V	0.9221 (0.2984)	0.0015 (0.0591)	0.0040 (0.0174)	0.0100 (0.0050)	1.0064
σ_ε	0.1184 (0.0201)	-	-	0.0911 (0.0124)	0.1720
K	0.0079 (0.0365)	-	-	0.2490 (0.0107)	0.1271
N : 888. Simulated Maximum Likelihood: -10,969.6					

Table 11: Importance Sampling parameter estimates. Importance Sampling parameter estimates of parameters for lognormal parameterization when we do not allow non-entrants to have paid the entry cost. Standard errors are in parentheses.

less than the reserve price. The main change is that the standard deviation of the distribution of $\sigma_{\varepsilon,a}$ increases quite significantly so that now the median value of $\sigma_{\varepsilon,a}$ across auctions is 0.4970. However, although this implies signals are less informative, there is still significant selection since the difference in the expected values of the marginal and inframarginal mill entrant are quite similar to the previous estimates (see Table 8). For both types of firms the reduced precision of the signal, and low K , leads them to use lower s 's as entry costs are low.

B Nested Pseudo Likelihood

The final estimator we consider is the Nested Pseudo-Likelihood estimator of Aguirregabiria and Mira (2007). We use this estimator to see how the results change when we relax the assumption that the parameters $\sigma_V^2, \sigma_\varepsilon^2$ and K are common across types within an auction and that the equilibrium played has mills (the type with higher mean values) using a lower s' . The NPL estimator, which to our knowledge has not previously been used in the auction literature, works by iterating two steps. In the first step, the structural parameters are estimated using the likelihood of firms' entry decisions and bids, given a set of beliefs about other players' strategies. In the auction context, these are beliefs about the distribution of the highest value of other entering firms, since it is this distribution that determines the probability that a firm wins and its surplus if it wins. In the second step, these beliefs are updated using the strategies implied by the parameters from the first step. If this process converges (which is not guaranteed in the case of the game), players will be playing best responses given their beliefs and beliefs will be consistent with players' strategies, so that strategies and beliefs will constitute a Bayesian Nash equilibrium. The appendix gives more complete details.

The NPL estimator is potentially consistent for any specification of initial beliefs. However, if the process is started from beliefs that are close to potential entrants' equilibrium beliefs, then it can be hoped - although it is not guaranteed - that the procedure will converge to the particular equilibrium that is being played, even if it does not satisfy the restrictions that we make in our other procedures. To get our initial estimates of beliefs we estimate a parametric distribution (Weibull) for the value of the highest bid, having controlled for observable covariates in a first-step.³³

Here we describe the nested pseudo-likelihood procedure in greater detail since we are not aware of it being used elsewhere in the empirical auction literature. Aguirregabiria and Mira (2002) and Aguirregabiria and Mira (2007) show how this procedure can be used to estimate both single agent models and games of incomplete information with a lower computational burden than nested fixed point methods. Applied to games, the procedure involves the iteration of two steps: in the first step, a pseudo-likelihood for each player's action is maximized to give estimates of the parameters, based on each player's best response to a set of beliefs about the actions of other players. In the first step, these beliefs are treated as data. In the second step, the new parameter estimates are used to

³³Note that this cannot give us a consistent estimate of what we actually want: a player's beliefs about the highest value of other players. That is because this object is not observed in the data since not all players can be assumed to bid their values. Therefore we have to use Nested Pseudo-Likelihood, which does not require consistent first stage estimates, rather than a simpler two-step estimator.

	Constant	Log(Sale Value)	Log(Logging Costs)	Density	Species Herfindahl	Log(Volume)	Std. Deviation	Median Value of Parameter
μ_{mill}	-8.0103 (0.07753)	3.1557 (0.0585)	-1.2716 (0.1699)	-0.0027 (0.0187)	-0.0149 (0.1263)	-0.0055 (0.0424)	0.3554 (0.0346)	3.7715
$\mu_{difference}$	-1.4463 (1.4704)	-0.0004 (0.0133)	-0.0020 (0.0112)	0.0030 (0.0290)	-0.0111 (0.0696)	0.0151 (0.0948)	0.5758 (1.7636)	-1.3451
σ_V	1.0904 (0.0137)	-0.0163 (0.0028)	0.0008 (0.0035)	- (0.0035)	- (0.0005)	- (0.0091)	0.0013 (0.0049)	1.0014
σ_ε	0.0039 (0.0238)	- (0.0238)	- (0.0299)	0.0055 (0.0299)	-0.0088 (0.0407)	0.0020 (0.0143)	0.7192 (0.1752)	0.497
K	0.0213 (0.0735)	- (0.0735)	- (0.0269)	0.0032 (0.0269)	0.0059 (0.0359)	0.0056 (0.0091)	0.0642 (0.0049)	0.0885
N : 888. Simulated Maximum Likelihood: -10,582.1								

Table 12: Importance Sampling parameter estimates. Importance Sampling parameter estimates of parameters for lognormal parameterization when we do allow non-entrants to have paid the entry cost. Standard errors are in parentheses.

update the beliefs. The two-step process is iterated until both beliefs and the parameters converge, at which point players will be playing best responses given their beliefs and beliefs will be consistent with players' strategies (i.e., strategies and beliefs will constitute a Bayesian Nash Equilibrium). The method contrasts with a nested fixed point procedure where an equilibrium is found at each step, and it also extends to cases where there is a finite mixture of unobserved heterogeneity.

Throughout much of this section we simplify the model that we are trying to estimate for expositional purposes. One important simplification we make is that we observe data generated by a second price auction. Our methodology is easily extended to cases where the data generating process is an English auction or a more general open outcry auction where bidders may not bid up to their values. We develop these extensions more fully in Appendix B. Finally, although we ignore reserve prices here, it is straightforward to include them, and we do in our empirical application.

Before describing the details of the method, we outline some additional notation. We then present the estimation method with multiple (observed) types of potential entrants, but with no unobserved heterogeneity between auctions. Then we describe how the routine changes when there is unobserved heterogeneity between auction types, before making some comments about how multiple equilibria may affect estimation.

B.0.1 Notation

To keep the notation simple, we will assume that $\bar{\tau} = 2$ (two types of potential entrants) for each auction and that within a type, all bidders are symmetric and use symmetric entry strategies. The type of each potential entrant is observed to the econometrician. Initially, we will also assume that all T observed auctions are identical i.e., for each auction bidders values are drawn from the same distributions F_τ^V (pdf f_τ^V), where $F_\tau^V = N(\mu_\tau, \sigma_V^2)$ and that N_τ is the same across auctions. The distribution of the signal noise around the true value is also normal, i.e., $F^\varepsilon = N(0, \sigma_\varepsilon^2)$. We will also assume that σ_v^2 , σ_ε^2 and K are the same across bidder types, although it is straightforward to relax this assumption. The complete set of parameters to be estimated is $\theta \in \{\mu_1, \mu_2, \sigma_v^2, \sigma_\varepsilon^2, K\}$.

To explain the estimation procedure, some additional notation will be useful. When a type τ bidder receives a signal s , his posterior belief is that his true value is distributed $F'_\tau(\cdot|s, \theta) = N(\alpha\mu_\tau + (1-\alpha)s, \sigma_V'^2)$, where $\alpha = \frac{1/\sigma_V^2}{1/\sigma_V^2 + 1/\sigma_\varepsilon^2}$ and $\sigma_V' = \sqrt{\frac{1}{1/\sigma_V^2 + 1/\sigma_\varepsilon^2}}$. $G_\tau^{-i}(\cdot)$ denotes the beliefs that a type τ bidder has about the distribution of the highest bid that will be placed by *other* bidders (hence the $-i$) when he makes his entry decision. $g_\tau^{-i}(\cdot)$ is the associated pdf. $s'_\tau(G_\tau^{-i}, \theta)$ is the signal which makes a type τ bidder indifferent about entering given the parameters and his beliefs. $s'_\tau(G_\tau^{-i}, \theta)$ is defined implicitly by the following equation, because the marginal entrant's expected profits from entering must be equal to zero:

$$\int_0^\infty \left(vG_\tau^{-i}(v) - \int_0^v v'g_\tau^{-i}(v')dv' \right) f'_\tau(v|s'_\tau, \theta)dv - K = 0 \quad (2)$$

Given G_τ^{-i} , θ and f'_τ , s'_τ can be solved for efficiently using a standard non-linear solver.

For each potential entrant i , we observe either that he does not enter or that he enters and

submits a bid equal to b . Denote the observed action $a_{i\tau}$, where $a_{i\tau} = 0$ if there is no entry and $a_{i\tau} = b_{i\tau}$ otherwise. Given (G_τ^{-i}, θ) , the probability that a type τ potential entrant i does not enter is:

$$\Pr(a_{i\tau} = 0 | G_\tau^{-i}, \theta) = \int_0^\infty F^\varepsilon(s'_\tau(G_\tau^{-i}, \theta) - v | \sigma_\varepsilon^2) f_\tau^V(v) dv \quad (3)$$

The probability (pmf) that a type τ potential entrant i enters and bids $b > 0$ (his value) is:

$$\Pr(a_{i\tau} = b | G_\tau^{-i}, \theta) = (1 - F^\varepsilon(s'_\tau(G_\tau^{-i}, \theta) - b | \sigma_\varepsilon^2)) f_\tau^V(b | \theta) \quad (4)$$

It is possible to make less restrictive modeling assumptions by altering the action probabilities. For example, we can drop the assumption that we observe the winning bidder's bid, which would be the case in an English Button auction. Alternatively, we could loosen the restriction that we observe all bidders' values, as might be the case in an open outcry auction (see for example Haile and Tamer (2003)). These alternative action probabilities can be used instead to form the pseudo-likelihood used in our estimation. These options are given in Appendix B.

Before estimation begins, we specify initial guesses of $G_\tau^{-i}(\cdot)$ for each type of player. With no unobserved auction heterogeneity, it is straightforward to approximate these distributions from the data using either parametric or non-parametric techniques.

The iterative pseudo-likelihood procedure has two steps. For a particular iteration k , the two steps proceed as follows:

Step 1 (maximum pseudo-log likelihood estimation). In this step the parameters θ are estimated using the best response probabilities (3) and (4) given values for $G_\tau^{-i}(\cdot)$. Formally:

$$\hat{\theta}^k = \arg \max_{\theta} \sum_{t=1}^T \sum_{\tau=1,2} \sum_{i=1}^{N_\tau} \log \Pr(a_{i\tau t} | G_\tau^{-i,k-1}, \theta) \quad (5)$$

For each value of the parameters, the pseudo-likelihood is calculated by solving for $s'_\tau(G_\tau^{-i}, \theta)$ for each player type, and then calculating the probability of the observed action for each potential entrant.

Step 2 (update G). In this step the parameter values and the final values of $s'_\tau(G_\tau^{-i}, \theta)$ from Step 1 are used to update the G_τ^{-i} distributions. For a bidder of type 1, the probability that the highest bid of other bidders is less than some value x is

$$\begin{aligned} \widehat{G_1^{-i,k}}(x) &= \left(\int_0^\infty F^\varepsilon(s'_1 - v | \hat{\theta}^k) f_1^V(v | \hat{\theta}^k) dv + \int_0^x (1 - F^\varepsilon(s'_1 - v | \hat{\theta}^k)) f_1^V(v | \hat{\theta}^k) dv \right)^{N_1-1} \\ &\times \left(\int_0^\infty F^\varepsilon(s'_2 - v | \hat{\theta}^k) f_2^V(v | \hat{\theta}^k) dv + \int_0^x (1 - F^\varepsilon(s'_2 - v | \hat{\theta}^k)) f_2^V(v | \hat{\theta}^k) dv \right)^{N_2} \end{aligned} \quad (6)$$

Steps 1 and 2 are iterated until both G and θ converge.³⁴

In practice, it was found that the algorithm converged more quickly if G^{-i} was updated only

³⁴In practice we set the tolerance level to 1.0E-6.

σ_ε	N_1	N_2	$\hat{\mu}_1$	$\hat{\mu}_2$	$\hat{\sigma}_V$	$\hat{\sigma}_\varepsilon$	$\hat{K}$
5.00	1	5	209.4751 (0.393)	200.2048 (0.447)	24.8229 (0.212)	4.8584 (0.177)	10.1479 (0.196)
5.00	2	5	209.6216 (0.206)	200.2784 (0.373)	25.0191 (0.171)	5.0690 (0.266)	10.2776 (0.291)
5.00	3	5	209.7727 (0.220)	199.9918 (0.389)	25.0671 (0.169)	4.8339 (0.183)	10.2015 (0.259)
0.55	2	4	209.5822 (0.236)	200.3087 (0.277)	25.0717 (0.122)	0.6748 (0.125)	10.3577 (0.241)
55.00	2	4	209.5344 (0.220)	197.4346 (3.413)	25.1296 (0.214)	50.8864 (7.426)	9.7548 (0.646)

Table 13: Recovering Parameters, No Unobserved Heterogeneity. The table displays estimates of model parameters assuming the correct Signal model. Based on generating $T = 5000$ auctions where $\mu_1 = 210$, $\mu_2 = 200$, $\sigma_V = 25$, $K = 10$. The cases when $\sigma_\varepsilon = 0.55$ and 55 reflect the S and LS models, respectively. Standard deviations are in parentheses.

partially in Step 2, i.e., we use a convex combination of $G_1^{-i,k-1}(x)$ and $\widehat{G_1^{-i}}(x)$.³⁵

This description assumes that there is no observed heterogeneity between auctions. In practice, one could parameterize μ as a function of observed covariates, e.g., $\mu_1 = X\beta$ and $\mu_2 = X\beta + \gamma$, as is done elsewhere in the literature (see for example Athey, Levin, and Seira (forthcoming)) and could allow for a different number of potential entrants in different auctions. In this case, we alter the computational routine to use a separate G_τ^{-i} and $s'_\tau(G_\tau^{-i}, \theta)$ for each distinct set of auction covariates. In practice this increases computational demands and so we control for observed heterogeneity in a pre-step regression described in Section 5.

B.1 NPL Monte Carlos

In this section we present several Monte Carlo experiments to illustrate our estimation method. We begin with the case with asymmetric bidders. We consider two bidder types who differ in their mean value. Table 13 displays the results from several experiments. The first three rows illustrate how we can recover the parameters of interest as the number of potential bidders varies for a reasonable amount of variation in the signal. We believe that one benefit of the model is that it is general enough to approximately match the two polar entry cases used in the literature so far. The next two rows attempt to illustrate this. They approximate a data generating process akin to the S and LS models, respectively. In the fourth row $\sigma_\varepsilon = 0.55$ and so bidders almost perfectly know their value prior to paying the entry cost. In the final row we set $\sigma_\varepsilon = 55$ and so bidders have almost no information regarding their value prior to entry. In either case we can recover the underlying parameters quite well.

³⁵In practice we only update 20% each time, i.e. we place weight 0.2 on the new estimate of the distribution.

	(1)	(2)
μ_{mill}	3.704 (0.023)	3.6690 (0.024)
μ_{logger}	2.711 (0.018)	3.1013 (0.041)
$\sigma_{V,mill}$	0.921 (0.012)	0.950 (0.016)
$\sigma_{V,logger}$	same as mill	0.801 (0.026)
$\sigma_{\varepsilon,mill}$	0.0990 (0.0005)	0.0312 (0.0095)
$\sigma_{\varepsilon,logger}$	same as mill	0.3001 (0.008)
K_{mill}	0.0016 (0.0003)	0.0003 (0.0000)
K_{logger}	same as mill	0.2251 (0.0964)
N	888 auctions	888 auctions

Table 14: Nested Pseudo-Likelihood parameter estimates. Nested Pseudo-Likelihood estimates of parameters for lognormal parameterization. The two specifications differ on their restrictions regarding signal variance. The latter model allows it to differ across types due to differences in value dispersion and noise dispersion. Standard errors are in parentheses.

B.2 NPL Results

Table 14 shows the estimates from two Nested Pseudo-Likelihood specifications. The first column shows the results when the parameters σ_V , σ_ε and K are constrained to be the same for mills and loggers. In the second column these parameters are allowed to be different across the firm types. We can permit these differences since we no longer make assumptions about equilibrium selection that require the parameters to be the same. The main findings are that loggers are estimated to have values that are less dispersed than mills - a plausible finding since loggers are less tied to the particular specifications of their manufacturing facility - and that their K and σ_ε are somewhat higher, suggesting that signals are less precise. However, because σ_V is also lower, their posterior belief about their value is still quite precise and the degree of selection for loggers in Table 8 is quite similar to the other specifications.³⁶

C Alternative Data Generating Processes

Here we outline how we can adapt the estimation procedure to loosen our restrictive assumptions about the data generating process.

³⁶If values and signals were both normally distributed, the bidder's posterior belief about his value would have variance equal to $\frac{1}{1/\sigma_V^2 + 1/\sigma_\varepsilon^2}$, so an increase in σ_ε^2 and a decrease in σ_V^2 tend to offset.

C.1 English Button Auction

If we assume that the data generating process is the classic English Button Auction where an auctioneer continuously raises the price and bidders press a button signaling their willingness to participate, we can interpret bidders' bids as their dropout points, which are their values. Thus any bid we observe is a bidder's value. However, we do not observe the winning bidder's value since he wins when the bidder with the second highest value removes his finger from the button. Therefore, while the probabilities of a bidder not entering and submitting an observed bid are as in Equations 3 and 4, respectively, we now lose the information that observing the winner's bid provides. Instead, we incorporate the probability of the winning bidder having a value in excess of the second highest bid b_2 by:

$$\Pr(a_{i\tau} = \text{win} | G_{\tau}^{-i}, \theta) = \int_{b_2}^{\infty} (1 - F^{\varepsilon}(s'_{\tau}(G_{\tau}^{-i}, \theta) - v | \sigma_{\varepsilon}^2)) f_{\tau}^V(v) dv \quad (7)$$

C.2 General Open Outcry Auction

The English Button auction, as imagined by Milgrom and Weber (1982), is obviously a simplification of what often are open outcry auctions. In these auctions we may worry that submitted bids are not actually bidders' values. Instead, a very general set of assumptions one can make upon observing a set of bids from a set of potential bidders, where the second highest bid b_2 , are:

- A1 The winning bidder had a value greater than b_2 .
- A2 The losing bidder that bid the most had a value equal to b_2 .
- A3 All participating bidders had values less than b_2 but greater than the reserve price r .

Probability $\Pr(a_{i\tau} = \text{win} | G_{\tau}^{-i}, \theta)$ is as in Equation 7. The probability that the highest losing bidder had a value equal to b_2 is as in Equation 4. However, for those bidders for which we observe a losing bid, we now say:

$$\Pr(a_{i\tau} = b \in [r, b_2] | G_{\tau}^{-i}, \theta) = \int_r^{b_2} (1 - F^{\varepsilon}(s'_{\tau}(G_{\tau}^{-i}, \theta) - v | \sigma_{\varepsilon}^2)) f_{\tau}^V(v | \theta) \quad (8)$$

References

- ACKERBERG, D. (2009): “A New Use of Importance Sample to Reduce Computational Burden in Simulation Estimation,” *Quantitative Marketing and Economics*, 7, 343–376.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70(4), 1519–1543.
- (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75(1), 1–53.
- ATHEY, S., AND J. LEVIN (2001): “Information and Competition at U.S. Forest Service Timber Auctions,” *Journal of Political Economy*, 109(2), 375–417.
- (2006): “Set-Asides and Subsidies in Auctions,” *Working Paper*.
- ATHEY, S., J. LEVIN, AND E. SEIRA (forthcoming): “Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions,” *Quarterly Journal of Economics*.
- BAJARI, P., H. HONG, AND S. RYAN (forthcoming): “Identification and Estimation of a Discrete Game of Complete Information,” *Econometrica*.
- BAJARI, P., AND A. HORTAÇSU (2003): “The Winner’s Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions,” *RAND Journal of Economics*, 34(2), 329–355.
- BALDWIN, L., R. MARSHALL, AND J.-F. RICHARD (1997): “Bidder Collusion at Forest Service Timber Sales,” *Journal of Political Economy*, 105(4), 657–699.
- BHATTACHARYA, V., J. ROBERTS, AND A. SWEETING (2010): “Public Procurement with Selective Entry,” *Working Paper*.
- BRANNMAN, L., AND L. FROEB (2000): “Mergers, Cartels, Set-Asides, and Bidding Preferences in Asymmetric Oral Auctions,” *Review of Economics and Statistics*, 82(2), 283–290.
- BULOW, J., AND P. KLEMPERER (1996): “Auctions Versus Negotiations,” *American Economic Review*, 86(1), 180–194.
- EINAV, L., AND I. ESPONDA (2008): “Endogenous Participation and Local Market Power in Highway Procurement,” *Working Paper*.
- ELICKSON, P., AND S. MISRA (forthcoming): “Enriching Interactions: Incorporating Outcome Data into Static Discrete Games,” *Quantitative Marketing and Economics*.
- ERTAÇ, S., A. HORTAÇSU, AND J. ROBERTS (forthcoming): “Entry into Auctions: An Experimental Analysis,” *International Journal of Industrial Organization*.
- HAILE, P. (2001): “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, 91(3), 399–427.

- HAILE, P., H. HONG, AND M. SHUM (2006): “Nonparametric Tests for Common Values In First-Price Sealed-Bid Auctions,” *Working Paper*.
- HAILE, P., AND E. TAMER (2003): “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111(1), 1–51.
- HARTMANN, W. (2006): “Intertemporal Effects of Consumption and Their Implications for Demand Elasticity Estimates, Quantitative Marketing and Economics,” *Quantitative Marketing and Economics*, 4(4), 325–349.
- HARTMANN, W., AND H. NAIR (forthcoming): “Retail Competition and the Dynamics of Demand for Tied Goods,” *Marketing Science*.
- HECKMAN, J. (1976): “The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models,” *Annals of Economic and Social Measurement*, 5, 475–492.
- HENDRICKS, K., J. PINKSE, AND R. PORTER (2003): “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions,” *Review of Economic Studies*, 70(1), 115–145.
- HENDRICKS, K., AND R. PORTER (2007): *An Empirical Perspective on Auctions* vol. 3 of *Handbook of Industrial Organization*, chap. 32. North-Holland.
- HU, Y., D. MCADAMS, AND M. SHUM (2009): “Nonparametric Identification of Auction Models with Non-Separable Unobserved Heterogeneity,” *Working Paper*.
- HUBBARD, T., AND H. PAARSCH (2009): “Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation,” *International Journal of Industrial Organization*, 27, 1–14.
- KIRKEGAARD, R. (2006): “A Short Proof of the Bulow-Klemperer Auctions vs. Negotiations Result,” *Economic Theory*, 28, 449–452.
- KRASNOKUTSKAYA, E. (forthcoming): “Identification and Estimation in Highway Procurement Auctions Under Unobserved Auction Heterogeneity,” *Review of Economic Studies*.
- KRASNOKUTSKAYA, E., AND K. SEIM (forthcoming): “Bid Preference Programs and Participation in Highway Procurement Auctions,” *American Economic Review*.
- LEVIN, D., AND J. SMITH (1994): “Equilibrium in Auctions with Entry,” *American Economic Review*, 84, 585–599.
- LI, T., AND X. ZHENG (2009): “Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions,” *Review of Economic Studies*, 76, 1397–1429.

- MARMER, V., A. SHNEYEROV, AND P. XU (2010): “What Model for Entry in First-Price Auctions? A Nonparametric Approach,” *Working Paper*.
- MEAD, W., M. SCHNIEPP, AND R. WATSON (1983): “Competitive Bidding for U.S. Forest Service Timber in the Pacific Northwest 1963-83,” Washington, DC: U.S. Department of Agriculture, Forest Service.
- MILGROM, P., AND R. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089–1122.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6(1), 58–73.
- PAARSCH, H. (1997): “Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbia Timber Sales,” *Journal of Econometrics*, 78, 333–357.
- PALFREY, T., AND S. PEVNITSKAYA (2008): “Endogenous Entry and Self Selection in Private Value Auctions: An Experimental Study,” *Journal of Economic Behavior and Organizations*, 66, 731–747.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2010): “Sequential Estimation of Dynamic Discrete Games: A Comment,” *Econometrica*, 8(2), 833–842.
- RILEY, J., AND W. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 58–73.
- ROBERTS, J. (2009): “Unobserved Heterogeneity and Reserve Prices in Auctions,” *Working Paper*.
- ROBERTS, J., AND A. SWEETING (2010a): “The Real Effect of Potential Competition in Merger Analysis,” *Working Paper*.
- (2010b): “When Should Sellers Use Auctions?,” *Working Paper*.
- SAMUELSON, W. (1985): “Competitive Bidding with Entry Costs,” *Economic Letters*, 17, 53–57.
- SWEETING, A. (2009): “The Strategic Timing of Radio Commercials: An Empirical Analysis Using Multiple Equilibria,” *RAND Journal of Economics*, 40(4), 710–742.