Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions*

Susan Athey, Jonathan Levin and Enrique Seira†

July 2004

Abstract

We study entry and bidding patterns in sealed bid and open auctions with heterogeneous bidders. Using data from U.S. Forest Service timber auctions, we document a set of systematic effects of auction format: sealed bid auctions attract more small bidders, shifts the allocation towards these bidders, and can also generate higher revenue.

We propose a model, which extends the theory of private value auctions with heterogeneous bidders to capture participation decisions, that can account for these qualititive effects of auction format. We then calibrate the model using parameters estimated from the data and show that the model can explain the quantitative effects as well. Finally, we use the model to provide an assessment of bidder competitiveness, which has important consequences for auction choice.

*We thank Phil Haile, Guido Imbens, Richard Levin, Paul Milgrom, Rob Porter, and Ilya Segal for helpful suggestions, and Jerry Hausman for advice at the early stages of this project. Athey and Levin acknowledge the support of the National Science Foundation.

†Athey: Stanford University and NBER, athey@stanford.edu; Levin: Stanford University, jdlevin@stanford.edu; Seira: Stanford University, eseira@yahoo.com.

1. Introduction

Auction design has become increasingly important in many markets. A central, and frequently debated, design issue concerns the relative performance of open and sealed bid auctions. This choice comes up in structuring sales of natural resources, art and real estate, in auctioning construction and procurement contracts, in asset liquidation sales, and in designing bidding markets for commodities such as electricity.

Economic theory provides on the one hand very little and on the other hand perhaps too much guidance on the merits of open and sealed bid auctions. The seminal result in auction theory, Vickrey's (1961) Revenue Equivalence Theorem, states that under certain conditions, the two formats have essentially equivalent equilibrium outcomes. Specifically, if bidders are risk-neutral and have independent and identically distributed values, the two auctions yield the same winner, the same expected revenue, and even the same bidder participation. In practice, however, these assumptions often seem too strong. Further work points out that as they are relaxed, auction choice becomes relevant, with the comparison between open and sealed bidding depending on both the details of the market (e.g. bidder heterogeneity, collusion, common rather than private values, risk-aversion, transaction costs and so on) and the designer's objective (e.g. revenue maximization or efficiency).

There has been less progress in providing empirical evidence on the performance of alternative auction designs. A difficulty is that many real-world auction markets tend to operate under a given set of rules rather than systematically experimenting with alternative designs. In this paper, we combine theory and empirical analysis to study the use of open and sealed bid auctions to sell timber from the national forests. The U.S. Forest Service timber program provides an excellent test case in market design as it has historically used both open and sealed auctions, at times even randomizing the choice. The timber sale program is also economically interesting in its own right. Timber logging and milling is a $100 billion a year industry in the U.S.,[1] and about 30% of timberland is publicly owned. During the time period we study, the federal government sold about a billion dollars of timber a year.

A long-standing debate surrounds the design of federal timber auctions. An early study by Mead (1967) argued that open auctions generated less revenue. In 1976,

[1] This number is from the U.S. Census and combines forestry and logging, sawmills, and pulp and paperboard mills (NAICS categories 113, 3221 and 321113).

Congress proposed the use of sealed bidding. The implementation of the law, however, allowed forest managers to use open auctions if they could justify the choice. As a result, sale method has varied geographically. In the Pacific Northwest, the largest Forest Service region, open auctions have predominated apart from a short period following the 1976 law. We focus instead on the neighboring Northern region comprised of Idaho and Montana, and provide additional evidence from California sales; both areas used a mix of formats during our sample period, 1982-1990.

The theoretical component of our analysis begins by highlighting two departures from the standard independent private value auction model, departures that are especially salient for timber auctions. First, we allow bidders to have heterogeneous value distributions. Here, we are motivated by the substantial variation among participants in Forest Service auctions, where the bidders range from large vertically integrated forest products conglomerates to individually-owned logging companies. Our second departure from the standard model is to endogenize participation by making it costly to acquire information and bid in the auction. Explicitly modeling participation decisions by heterogeneous bidders makes the model more realistic, and more importantly gives rise to new testable hypotheses about entry patterns.[2]

In our baseline model, we assume that entry decisions and bidding decisions are made independently by the firms, so that there is no collusion. This competitive theory generates predictions about how entry, allocation and revenue will vary by auction format. When bidders are homogeneous, auction format will have no systematic effect on any of these outcome variables. With heterogeneous bidders, however, sealed bidding promotes entry by bidders who are relatively weak from an ex ante standpoint and discourages entry by bidders who are relatively strong. The allocation in a sealed auction also shifts toward weaker bidders, both because of entry and because in the sealed bid equilibrium, weaker firms submit bids with a smaller profit margin than stronger firms. To see the logic for these findings, observe that with an open auction, the entrant with the highest value always wins. This makes weak bidders hesitant to spend money to participate if strong bidders are also likely to be present. In contrast, in a sealed bid auction, strong bidders have a relatively

[2]Maskin and Riley (2000) provide the seminal analysis of asymmetric first-price auctions with fixed participation. Several papers study entry decisions in auctions with symmetric bidders, but discussion of entry with asymmetric bidders has been limited to examples. Milgrom (2004, chapter 6) provides an insightful overview.

large incentive to shade their bids below their true valuations, so a weaker bidder can win despite not having the highest valuation. This handicapping effect promotes the entry of weaker bidders and may discourage the entry of strong bidders.[3]

We also argue that when entry costs are similar across bidders, the entry patterns of weak bidders will be more sensitive to auction format than the entry patterns of strong bidders. Because strong bidders expect higher profits than weak bidders, weak and strong bidders with similar entry costs cannot both be indifferent about entry. If the set of potential strong bidders is fixed for a given auction (e.g. the set of local mills in a timber auction), then these bidders may enter regardless of format, while weak bidders will be marginal entrants whose participation is sensitive to the auction format.

The competitive theory does not generate unambiguous predictions about revenue. The existing examples and numerical results suggest that with a fixed set of heterogeneous bidders, revenue is typically (though not always) higher with sealed bidding. Endogenous entry, however, generates an additional complication because participation varies with the auction format. A revenue comparison, therefore, will depend on all the primitives of the model: the value distributions of the bidders together with entry costs. Consequently one of our goals in this paper is to estimate these primitives in order to compare the revenue gain (if any) from sealed bidding to the efficiency distortion that sealed bidding induces in both entry and bidding.

Our empirical investigation of timber auctions has two parts. In the first part, we test the qualitative predictions of our theory, and quantify the effect of auction format on observed outcomes. The second part of the analysis exploits an additional assumption about behavior, namely that in sealed bid auctions, bidders behave according to our competitive theory. We estimate the primitives of the model, and use our estimates to assess whether the theory can account for the quantitative differences across formats we observe in the data, as well as to quantify the trade-offs in revenue

[3]The first half of this argument figured prominently in the design of the British 3G spectrum auction. Initially, the proposed number of licenses for sale was to be just equal to the number of incumbent firms, whose entry was assured. Concerned that there would be little de novo entry to raise prices in an open auction, the lead designer Paul Klemperer proposed an "Anglo-Dutch" design that would have added a sealed bid component to the auction. Ultimately concerns about entry were alleviated by adding another license and it was possible to run a successful open auction. Interestingly, the Netherlands ran an open auction for 3G spectrum despite having a number of licences just equal to the number of incumbents. This resulted in minimal entry and very low prices.

and efficiency suggested by theory.

For both parts of the empirical work, we classify the bidders into two groups: mills that have manufacturing capacity and loggers that do not. We provide a variety of evidence that mills tend to have higher values for a given contract than logging companies, which have to re-sell the timber if they win the contract.

Our first test of the theory concerns bidder entry. We find that, conditional on sale characteristics, sealed bidding induces significantly more participation by loggers. Mill entry is roughly the same across auction formats in the Northern forests, and somewhat lower in the sealed bid auctions in California. We find evidence that sealed bid auctions are more likely to be won by loggers; this effect is substantial in the California forests and smaller (and only marginally significant) in the Northern forests. Finally, we measure winning bids to be 12-18% higher in the sealed bid auctions in the Northern forests. In the California forests, the difference is small and cannot be statistically distinguished from zero.

Although the theoretical model is qualitatively consistent with these results, it is less clear whether the quantitative differences in the auction outcomes can be accounted for by the competitive theory. In particular, the question arises of whether the competitive bidding model can reconcile both the large revenue gap in the Northern forests, and the minimal revenue effect in California.

To address this, it is useful to articulate alternatives and extensions to the baseline competitive model that might be able to rationalize our findings. We argue that several factors that seem plausible in the context of timber auctions, but are omitted from our baseline model, such as common values and bidder risk-aversion, are not good candidates. Instead, we focus on the possibility that behavior is not fully competitive in open auctions. Bidder collusion has been a long-standing concern in timber auctions; the prevailing view is that open auctions are more prone to collusion because bidders are face-to-face and can respond immediately to opponents' behavior. For this reason, we extend the theory to allow for collusion among strong bidders at open auctions. We show that collusion at open auctions need not affect the model's predictions for entry and allocation, but increases the predicted revenue difference between auction formats.

In the final part of the paper, we turn to a quantitative assessment of the alternative theories, focusing on the Northern forests. We use the techniques pioneered by

Guerre, Perrigne and Vuong (2001) to recover the distributions of bidder values from the sealed bidding data, under the assumption that observed bids are set to maximize profits against the empirical bid distribution. We also estimate the distribution of logger entry in sealed bid auctions, and combine this with the profits implied by the estimated value distributions to construct estimates of entry costs.

We use these estimates to make (out-of-sample) predictions about what would happen in open auctions under alternative behavioral assumptions, and compare these predictions to the actual open auction outcomes. This allows us to consider several questions including whether the theoretical model can explain the departures from revenue equivalence observed in the data, whether open auction behavior seems more consistent with competitive bidding or a degree of collusion, and whether bidder competitiveness might differ across regions.

Our results suggest that the estimated model can do plausible job of explaining both the differences in participation and the differences in allocation we observe across formats. Focusing on prices, we find that neither the assumption of perfectly competitive behavior, nor an assumption that mills collude perfectly at open auctions, can match the observed open auction prices. Rather, the data appears consistent with a mild degree of cooperative behavior on the part of participating mills.

Turning to the welfare differences between open and sealed bid auctions, we find that for a fixed set of participants, our calibrated model predicts relatively small discrepancies between sealed bid auctions and competitive open auctions. Sealed bid auctions raise more revenue, and distort the allocation away from efficiency and in favor of loggers, but the effects are small (less than 1%). The differences are somewhat larger when we account for equilibrium entry behavior: we predict that sealed bidding increase revenues by roughly 4-6% relative to a competitive open auction, at the cost of about 1% of social surplus.

We also observe that even a mild degree of collusion by the mills at open auctions — the behavioral assumption most consistent with the observed outcomes in the Northern forests — results in much more substantial revenue differences (on the order of 10-20%). This suggests that bidder competitiveness merits considerable attention in the choice of auction format. We conclude by discussing some preliminary estimates from the California forests. There we find that the observed outcomes are more consistent with competitive behavior in open auctions and, as a result, there is no dominant factor in the welfare comparison between open and sealed bid auctions.

Our paper is the first empirical study we are aware of that focuses on differential entry and the importance of bidder heterogeneity across auction formats.[4] Several prior studies have looked directly at revenue differences between open and sealed bid timber auctions. Johnson (1979) and Hansen (1986) study sales in the Pacific Northwest following the passage of the 1976 sealed bidding mandate. They reach conflicting conclusions: Johnson finds that the sealed bid auctions raised more revenue, while Hansen argues that the differences are insignificant after accurately accounting for sale characteristics. The episode is not, however, an ideal testing ground. As Hansen points out, the choice of auction format during this period was sensitive to lobbying, creating a potentially severe endogeneity problem that is hard to address empirically. Moreover, one might naturally be skeptical of testing equilibrium predictions in an unexpected and transient episode.

Subsequently, Shuster and Nicolluci (1993) and Stone and Rideout (1997) looked, respectively, at sales in Idaho and Montana and in Colorado. Both papers find higher revenue from sealed bid auctions. A nice feature of Shuster and Nicolluci's paper is that they exploit the often-random assignment of auction format in some of the Northern forests. Though we address a broader set of questions and from a somewhat different perspective, we have drawn on their work to select our data sample.

Our work is also related to the empirical literature on collusion at auctions. A variety of approaches have been suggested to assess whether bidding data are consistent with models of competition or collusion.[5] Some approaches require prior knowledge about the existence and structure of a cartel, while others interpret departures from symmetric bidding behavior as evidence of collusion. Our method differs in that we use behavior under one set of auction rules (sealed bidding) as a benchmark from which to evaluate the competitiveness of behavior under an alternative set of rules.

[4]Indeed, most analyses of auctions assume that bidders are symmetric. A few notable exceptions study asymmetries in auctions with fixed participation, including Bajari (1997), Brannman and Froeb (2000), Pesendorfer (2000), Jofre-Benet and Pesendorfer (2003), and Brendstrup and Paarsch (2003a, 2003b).

[5]Examples include Porter and Zona (1993, 1999), Bajari and Ye (2003), Pesendorfer (2000); see Bajari and Summers (2002) for a survey. Baldwin, Marshall, and Richard (1997) also analyze collusion in U.S. Forest Service timber auctions using data from open auctions, and they argue that collusion provides a better fit than competition.

2. Comparing Auctions: Theory

This section develops the theoretical model we use to frame our empirical analysis. Our starting point is the heterogeneous private values setting studied by Maskin and Riley (2000). With an eye toward the empirical patterns outlined above, we expand their analysis to make participation endogenous and to incorporate possible collusion in open auctions. In this exercise, there are numerous specific modeling choices to be made. To ease exposition, we begin with a baseline model, then discuss how the results change under alternative assumptions.

A. The Model

We consider an auction for a single tract of timber. Prior to the sale, the seller announces a reserve price r and the auction format: open ascending or first price sealed bid. There are N potential risk-neutral bidders. Each bidder i has a private cost k_i of gathering information and entering the auction. By paying k_i, bidder i learns his (private) value for the tract, v_i, and may bid in the auction. We refer to bidders who acquire information as *participants*, and denote the set of participants by n.

Entry costs and values are assumed to be independent across bidders. We model entry costs as draws from a common distribution $H(\cdot)$ with support $[\underline{k}, \overline{k}]$, and each bidder i's value as a draw from a distribution F_i with support $[\underline{v} = r, \overline{v}_i]$.[6] Anticipating our empirical analysis, we allow for two kinds of bidders. Bidders $1,...,L$ are *Loggers* and have value distribution F_L, while bidders $L+1,...,N$ are *Mills* and have value distribution F_M. We assume that F_M stochastically dominates F_L, so we sometimes refer to the mills as strong bidders and the loggers as weak bidders.

Assumption *(i)* F_L,F_M have continuous densities f_L,f_M; and *(ii)* for all v, $\frac{f_M(v)}{F_M(v)} \geq \frac{f_L(v)}{F_L(v)}$.

We adopt a standard model of the bidding process. In an open auction, the price rises from the reserve price and the auction terminates when all but one participating bidder has dropped out. With sealed bidding, participating bidders independently

[6] The assumption that the reserve price equals the lowest possible value is easily relaxed.

submit bids; the highest bidder wins and pays his bid. For both auctions, we assume that bidders make independent decisions to acquire information, but learn the identities of other participants before submitting their bids.[7]

A strategy for bidder i consists of a *bidding strategy* and an *entry strategy.* A bidding strategy $b_i(\cdot;n)$ specifies i's bid (or drop-out point in the case of an open auction) as a function of his value and the set of participating bidders. An entry strategy specifies whether he should participate as a function of his entry cost. An optimal entry strategy is a threshold rule, with bidder i entering if and only if his cost lies below some threshold K_i.

A *type-symmetric entry equilibrium* is a pair of bidding strategies $b_L(\cdot;n)$,$b_M(\cdot;n)$ and entry cost thresholds K_L,K_M with the property that: *(i)* loggers use the strategy b_L,K_L and mills the strategy b_M,K_M; *(ii)* each bidder's bid strategy maximizes his profits conditional on entering; and *(iii)* each bidder finds it optimal to enter if and only if his entry cost lies below his cost threshold. As is often the case with entry models, there may be many equilibria; as a result, our results compare sets of equilibria across auction methods.

B. Sealed Bid Auctions

We analyze the sealed bid auction in two steps. We first characterize optimal bidding for an arbitrary set of participants. We then characterize equilibrium entry. To focus on the main ideas, we defer proofs to the Appendix.

Suppose i is a participating bidder with value v_i. His expected profit is:

$$\Pi_i^s(v_i;n) := \max_{b\geq r}(v_i - b) \prod_{j\in n\setminus i} G_j(b;n), \tag{1}$$

where $G_j(b;n)$ is the probability that j will bid less than b. In equilibrium, bid strategies will be continuous and strictly increasing, so $G_j(b;n) = F_j(b_j^{-1}(b;n))$.

The first order condition for i's bidding problem is:

$$\frac{1}{v_i - b_i} = \sum_{j\in n\setminus i} \frac{g_j(b_i;n)}{G_j(b_i;n)}. \tag{2}$$

[7] This assumption is not essential. Indeed an earlier version of the paper assumed bids were submitted without information about opponent's participation. There we showed the same results under a modification of Assumption *(ii)*.

The first order conditions, together with the boundary condition that $b_i(r;n) = r$ for all i, uniquely characterize optimal bidding strategies (Maskin and Riley, 2000). These bid strategies are type-symmetric.

To identify the equilibrium entry thresholds, observe that a bidder should enter whenever his expected profit exceeds his entry cost. Given a set of entry thresholds $K = (K_1, ..., K_n)$, bidder i's expected profit from entry is:

$$\Pi_i^s(K) = \sum_{n \subset N} \left\{ \int \Pi_i^s(v_i; n) dF_i(v_i) \right\} \Pr\left[n \mid K, i \in n\right], \tag{3}$$

where $\Pr\left[n \mid K, i \in n\right]$ is the probability that the set of participants will be n given that i enters and opponents use their specified entry strategies.

The equilibrium entry cost thresholds satisfy:

$$K_i^s = \min\{\Pi_i^s(K^s), \overline{k}\}. \tag{4}$$

Proposition 1 *A type-symmetric entry equilibrium exists in the sealed bid auction. In equilibrium: (i) mills submit higher bids: $G_M(b;n) \leq G_L(b;n)$ for all $b \geq r$, despite the fact that (ii) mills shade their bids more than loggers bidders: $b_M(v;n) \leq b_L(v;n)$ for all v.*

The first part of the Proposition states that mills will tend to submit higher bids than loggers. From an empirical standpoint, this will provide a straightforward test of whether we have accurately classified mills as stronger than loggers. The second part of the Proposition states that mills shade their bids more loggers, a natural result given that the mills face weaker competition. The consequence is that a logger may win despite not having the highest value. We will show that, relative to an open auction, this provides an extra incentive for loggers to participate.

C. Open Auctions

We now turn to the open auction. We initially consider the case where behavior is competitive and discuss collusion below.

In an open auction, it is a dominant strategy for each participant to bid until the price reaches his valuation. Therefore $b_i(v;n) = v$ for all bidders i. Bidder i's expected profit, conditional on entering and having value v_i is:

$$\Pi_i^o(v_i;n) := \max_{b\geq r}(v_i - \mathbb{E}[\max\{b_{-i},r\}|b_{-i} \leq b]) \prod_{j\in n\setminus i} G_j(b;n), \tag{5}$$

where b_{-i} are the competing bids and $G_j(b;n)$ is the probability that j bids less than b. In equilibrium, $G_j(b;n) = F_j(b)$.

We identify equilibrium entry just as in the sealed bid case. Bidder i's expected profit as a function of the entry cost thresholds is:

$$\Pi_i^o(K) = \sum_{n\subset N} \left\{ \int \Pi_i^o(v_i;n)dF_i(v_i) \right\} \Pr\left[n \mid K, i \in n\right].$$

In equilibrium, each bidder enters if his expected profit exceeds his entry cost. So the equilibrium entry cost thresholds satisfy:

$$K_i^o = \max\{\Pi_i^o(K^o), \overline{k}\}. \tag{6}$$

Proposition 2 *A type-symmetric entry equilibrium exists in the open bid auction. In any such equilibrium, (i) mills submit higher bids:* $G_M(b;n) \leq G_L(b;n)$ *for all* $b \geq r$*, and (ii) all entrants bid their true value,* $b_i(v;n) = v$ *for all* i,v*.*

In equilibrium, mills enter more often and bid more conditional on entering. Moreover, the open auction is efficient in the sense that the participant with the highest value always wins. As we will see, this tends to discourage the entry of weaker bidders relative to the sealed bid case.

D. Comparing Auction Formats

We now present our main comparative results. As a point of reference, we start with the case where the bidders have identical value distributions. Here, an extension of the revenue equivalence theorem implies that so long as we restrict attention to symmetric equilibria, the open and sealed bid auctions have equivalent outcomes.

Proposition 3 *(Revenue Equivalence) If bidders are homogenous, so* $F_L = F_M$*, the sealed bid and open auction each have a unique symmetric entry equilibrium, in which* *the highest valued entrant wins the auction. These equilibria have (a) the same expected entry, and (b) the same expected revenue.*

Revenue equivalence breaks down if bidders are not homogenous. To analyze this case, we exploit the relationship between a bidder's equilibrium profits and his probability of winning. Given a value i and a set of participants n, bidder i's expected profit:

$$\Pi_i(v;n) = \int_{\underline{v}}^{v} \Pr[i \text{ wins} \mid v_i = x; n]dx. \tag{7}$$

This representation holds for both auction formats; it follows from applying the envelope theorem to the optimization problems (1) and (5).

We saw above that in a sealed bid auction with heterogeneous bidders, mills shade their bids more than loggers, while all bidders use the same strategy in an open auction. Therefore for any given set of opponents, a logger has a greater chance to win a sealed auction and hence higher expected profits. The argument is reversed for mills, leading to the following result.

Proposition 4 *For any type-symmetric entry equilibrium of the sealed bid auction, there is a type-symmetric entry equilibrium of the open auction in which: (1) Loggers are less likely to enter; (2) Mills are more likely to enter; (3) It is less likely a logger will win.*[8]

Because the sealed bidding equilibrium distorts the allocation toward loggers, only the open auction is efficient given a set of participating bidders. The next Proposition states that the efficiency of the open auction extends to entry.

Proposition 5 *(Efficiency) The socially efficient type-symmetric entry profile is an entry equilibrium of the open auction. Every sealed auction equilibrium is inefficient.*

As noted in the Introduction, there is unfortunately no general theoretical comparison for expected revenue (Maskin and Riley, 2000). Existing numerical examples suggest that, with participation fixed, sealed bid auctions often (but not always) result in higher revenue (Li and Riley, 1999). In terms of revenue, endogenous entry

[8]The statement of the result is complicated slightly by the fact that there may be several (type-symmetric monotone) entry equilibria for each auction format. If both formats have a unique entry equilibrium, loggers necessarily enter and win more with a sealed format.

could in principle either tip the revenue comparison further toward sealed bidding (if the primary entry effect is on loggers) or toward open bidding (if the primary entry effect is on mills). Therefore a revenue comparison demands a carefully parameterized model, which we develop in Section 5.

E. Collusion in Open Auctions

Collusion in open auctions has been a long-standing concern in Forest Service timber auctions (U.S. Congress, 1976; Froeb and McAfee, 1988; Baldwin et al, 1997). Here we consider the possibility of collusion by the mills in open auctions.

As collusive schemes can take many forms, we assume for concreteness that participating mills at an open auction are able to collude perfectly, so the participating mill with the highest value bids his value, while the other mills register as participants but do not actively bid).[9] Loggers simply bid up to their value. We maintain the assumption that bidders make independent participation decisions. In making their decisions, therefore, mills anticipate colluding with other participating mills, but do not coordinate entry.[10]

Fixing the set of participants, collusion clearly will lower revenue and increase mill profits. It has no effect on who wins the auction or on logger profits, because the high-valued mill is the relevant competitor for loggers in any case. Therefore, relative to the case of competition, mills have a greater incentive to participate; this in turn crowds out logger participation.

Proposition 6 *(Collusion) For any type-symmetric entry equilibrium of the open auction, there is a type-symmetric collusive equilibrium in which: (1) Loggers are less likely to enter; (3) Mills are more likely to enter; (3) It is less likely a logger will win. Thus, for any type-symmetric entry equilibrium of the sealed bid auction, there is a type-symmetric collusive equilibrium of the open auction where (1)-(3) hold.*

[9] More generally, to avoid arousing suspicion, the mills with lower values might place bids in the open auction at a point where many bidders are still active and thus the bids are unlikely to determine the auction outcome.

[10] There are forms of collusion, of course, that involve coordinated entry (such as bid rotation). We have looked for evidence of this in our data by checking whether the entry of pairs of mills or loggers is negatively correlated conditional on sale characteristics. There are a handful of pairs for which entry is significantly negatively correlated, but this pattern does not appear to be a strong feature of the data.

An important point is that relative to equilibrium outcomes of the sealed bid auction, the competitive and collusive outcomes of the open auction look qualitatively similar (lower prices, less logger entry, fewer sales won by loggers). The difference is one of magnitude.

F. Discussion of Modeling Choices

In this section, we briefly discuss a few of our modeling choices. We first discuss our model of entry. We then consider two issues omitted from the model: common values and bidder risk-aversion.

Concentrated versus Dispersed Entry Costs

Our model assumes that bidders differ in their costs of information acquisition and bid preparation.[11] In principle these differences could be either large or small (relative to the average entry cost); this distinction turns out to be relevant in interpreting the results.

Consider first that entry costs are dispersed. In this case, every potential bidder will be "marginal" in the sense of having a probability of entry strictly between zero and one. Moreover, a change in the auction format (which changes all bidders' expected profits) will affect the equilibrium entry behavior of all bidders — both mills and loggers.

In contrast, if entry costs are concentrated, it cannot be true that both mills and loggers are marginal, because mills expect higher profits than loggers. In equilibrium, either mills will be roughly indifferent to entering while loggers expect strictly negative profits and don't enter (clearly not the appropriate assumption for our data), or alternatively, loggers will be roughly indifferent while mills always enter.[12] In the latter case, mill participation will be unaffected by auction format, while logger participation will be strictly higher with sealed bidding. An interesting consequence is the effect of sealed bidding on revenue via its effect on participation will always be positive.

[11] These differences might correspond to knowledge of an area, or the availability of personnel.

[12] In the former case, the relevant bidders are homogenous so revenue equivalence holds across auctions. A third possibility is that all bidders agree on whether or not entry is profitable. In this case, the set of participating bidders is effectively fixed in a given auction. A fourth (and somewhat perverse) possibility is that all loggers enter, and given this, mills strictly prefer not to enter. We disregard this equilibrium.

Common Values and Risk-Aversion

In timber auctions, differences in bidder costs and contractual arrangements provide a source of private value differences. At the same time, bidders can obtain private estimates of the quality and quantity of timber, which suggests a potential "common value" component as well (Athey and Levin, 2001).[13] Haile (2001) studies how resale markets in timber auctions can lead to common values even if the underlying environment has private values. In the presence of common values, expected revenue is higher in open auctions, at least with symmetric bidders.

Bidder risk-aversion also has implications for the comparison between open and sealed bid auctions (see e.g. Matthews, 1987). If bidders are symmetric and have CARA or DARA preferences, expected revenue is higher with a sealed bid auction, while participation is higher at open auctions. It is plausible that bidders at Forest Service timber auctions might exhibit risk-aversion; Athey and Levin (2001) provide some indirect support for this based on the way observed bids are constructed.[14]

Without dismissing the possibility of either common values or bidder risk-aversion, we decided not to focus on them in our theoretical model for two reasons. First, incorporating either greatly complicates the analysis. Second, to jump ahead, our empirical results suggest that neither common values nor risk-aversion are the primary cause of the departures we observe from revenue equivalence.

3. Timber Sales

The U.S. Forest Service has historically used both open and sealed bid auctions to sell timber from the national forests. In this section, we describe the mechanics of a timber sale, the data for our study, factors that relate to the auction format, and how we classify competing bidders.

A. The Timber Sale Process

[13] Athey and Levin (2001) show that in certain Forest Service auctions, bidders can profit from acquiring commonly relevant information about timber volumes. They also show, however, that the potential rents are competed away, suggesting that the equilibrium information asymmetry about volumes may not be quantitatively large.

[14] See also Perrigne (2003) for evidence of risk aversion from French timber auctions.

Our data consists of timber sales held between 1982 and 1990 in Lolo and Idaho Panhandle National Forests, neighboring forests on the Idaho/Montana border. These are the two forests in the Forest Service's Northern region with the largest timber sale programs. They make a good test case for comparing auction formats because they use a mix of open and sealed auctions and the tracts sold under the two formats appear to be relatively homogenous. We discuss the way auction format is determined in more detail below. In Section 4C, we provide additional evidence from forests in the Pacific Southwest region. These California forests also use both open and sealed bidding, but the auction format varies more systematically with the size of the sale, which makes controlling for tract differences more challenging.

In both regions, a sale begins with the Forest Service identifying a tract of timber to be offered and organizing a "cruise" to estimate the merchantable timber. The sale is announced publicly at least thirty days prior to the auction. The announcement includes estimates of available timber and logging costs, tract characteristics and a reserve price. It also states whether the auction will involve open or sealed bids. In some cases, the Forest Service restricts entry to firms with less than 500 employees. We do not consider these small business sales — in principle the bidders are more homogenous than in regular sales, removing what we believe to be a crucial factor in distinguishing open and sealed sales.[15]

Following the sale announcement, the bidders have the opportunity to cruise the tract and prepare bids. As in the model, we classify bidders into two types: mills that have manufacturing capability and logging companies that do not. We discuss this classification below.

After the auction is completed, the winner has a set amount of time – up to seven years but more often one to four years in our sample — to harvest the timber. Some of the sales in our sample are "scale sales" meaning the winner pays for the timber only after it is removed from the tract. The fact that payments are based on harvested timber, but bids are computed based on quantity estimates means there can be a gap between the winning bid and the ultimate revenue. Athey and Levin (2001) study the incentive this creates for strategic bidder behavior. For the scale sales in our sample, we have limited harvest data, so we use the bid price as a proxy for revenue. The

[15]Ideally we would be able to perform a separate comparison of open and sealed bidding using the small business sales. Unfortunately, we do not observe enough to make a good comparison.

remaining sales are “lump-sum” sales. In these sales the winner of the auction pays the bid price directly.

B. Data Description

For each sale in our sample, we know the identity and bid of each participating bidder, as well as detailed sale characteristics from the Forest Service appraisal. This is the same appraisal information provided to bidders. Table 1 presents some basic summary statistics.

Focusing on the full sample, there are some obvious differences between the open and sealed bid auctions. The average sale price per unit of timber (in 1983 dollars per thousand board feet of timber or $/mbf) is roughly $70 in the open auctions and $80 in the sealed auctions. The number of entering loggers is also somewhat higher in sealed auctions (3.4 versus 2.6), while the number of entering mills is slightly lower (1.5 versus 1.2). Contracts sold by sealed auction are more likely to be won by a logging company than tracts sold by open auction.

These numbers are broadly consistent with the model presented above. At the same time, the Table indicates that the tracts sold by open auction are not identical to those sold by sealed bid. While the per-unit reserve price of the timber is similar across format, the open auction tracts tend to be larger. The average open auction has an estimated 2893 mbf of timber, while the average sealed bid sale has only 1502 mbf. This suggests that we need to understand how the sale format is decided and control for tract characteristics if to isolate the effects of auction format.

C. Choice of Sale Method

In Forest Service timber sales, the choice of sale method is made locally by forest managers. One reason for focusing on the two Northern forests is that Shuster and Nicolluci (1993) report that for a subset of these sales, the choice of sale format was explicitly randomized. In one forest district the format apparently was determined by picking colored marbles out of a bag. Unfortunately, we do not know precisely how the randomization procedure varied across forest districts and over time. We get similar empirical results focusing on the subset that Shuster and Nicolluci (1993) identify as randomized (though our estimates are somewhat less precise due to the smaller sample size).[16]

[16] Within our two forests, we include more districts and years than those Shuster and Nicolluci

To better understand the determinants of sale method in our sample, we consider a logit regression where the dependent variable is a dummy equal to 1 if the auction is sealed bid and equal to 0 if the sale is an open auction. We include a large set of observable tract characteristics, including the reserve price, the Forest Service estimates of the volume of timber, its eventual selling value, and the costs of logging, manufacturing and road-building. We also include the density of timber on the tract, the contract length, whether the sale is a salvage sale, and a Herfindal index of the concentration of species on the tract. To capture market conditions, we include the number of U.S. housing starts in the previous month. Finally, as measures of potential competition, we use the number of logging firms and sawmills in the county of the sale, as counted by the U.S. Census in the past year, as well as the number of "active" logging firms and sawmills, constructed as the number of distinct bidders in the same forest district in the 300 days preceding the sale.[17] We also include dummy variables for the year of the sale, the quarter of the sale, the forest district in which the sale took place and if major species were present. We are particularly sensitive to the importance of sale size, so rather than simply assuming a linear or quadratic effect, we specify its effect as a step function with 10 steps (that roughly correspond to deciles in the data).[18]

The results are reported in Table 2. As expected, sale size is a significant correlate of auction method. Even after controlling for time and geographic location, smaller sales tend to be sealed bid, while larger sales tend to be open auctions. Moreover,

identify as randomized (they focus on 1987-1990). In including these additional years, our motivation is that the set of tracts sold by open and sealed bidding appear to vary mainly with size, time and location, precisely the characteristics we need to control for in any case with the randomized sales. We focus on the two largest Northern forests because timber markets in Idaho and Montana are quite local due to the geography, while tract characteristics also vary with geography as well, making it difficult to effectively control for heterogeneity in forests with fewer sales.

[17]In terms of capturing potential competition, these measures probably suffer from a degree of measurement error. Apart from the fact that logging firms may go in and out of business without our knowledge, the Forest Service data records bidder names with a variety of spellings and abbreviations. Despite sale by sale checking of the names and cross-referencing with industry reference books, in the case of very small firms that appear relatively few times in the data it is sometimes hard to distinguish whether two bidders in distinct sales are really the same firm. This is less of problem with mills as their manufacturing capability is coded, there a far fewer in total, and they generally appear many times. Note that for the California sales, we use the forest rather than the district as a unit of analysis, because forests are smaller there.

[18]We use this functional form in all our regressions. We have also tried using a series expansion for volume and splines, with similar results throughout.

different forest districts use somewhat different sale methods on average.

Because sale method varies with observable sale characteristics, we want to control for these characteristics in comparing the outcomes of the open and sealed bid auctions. A concern is that, even controlling for tract characteristics flexibly, some open sales in our data may look very "unlike" any sealed bid sales and conversely some sealed sales may look unlike any open sales. This will be reflected in having some sales for which, conditional on characteristics, the predicted probability of being sealed or open according to our logit regression will be close to zero or one. Figure 1 plots a smoothed histogram of these predicted probabilities, also called the propensity score. As can be seen, there are some sales that are cause for concern. To alleviate this in our empirical analysis below, we drop sales that have a propensity score below 0.075 or above 0.925. This results in dropping 129 open auctions and 8 sealed auctions.[19]

A problem we cannot easily solve is that the choice of auction method may depend on characteristics of the sale observed by the bidders and the Forest Service, but not in our data. In this case, a regression of entry or revenue on auction method, even controlling for observed characteristics, will have an endogeneity problem. We discuss this possibility at more length in Section 4E.

D. Bidder Heterogeneity: Mills and Loggers

We try to capture the diversity of bidders by distinguishing between mills (formally, firms with manufacturing, which are larger and can process at least some of the timber themselves, and logging companies, who must re-sell all the timber they harvest. This distinction is just one of several we could draw, but in practice it turns out to be similar to other natural classifications. For instance, we have categorical data on firm employment and find that if we break the firms into large and small employers, we arrive at very nearly at the same classification.[20] Mills also attend more auctions than most loggers, although there are a few loggers who attend frequently.

Our theoretical model assumes that mills tend to have higher willingness to pay than loggers. The theory suggests several ways to check this assumption — mills

[19]The dropped sales are generally large volume sales in districts that ran few sealed auctions.

[20]The employment data appears to be somewhat noisy, but to convey a rough sense, suppose we classify bidders as "large" if they have more than XX employees. Then of the 1536 appearance by mills in our data, 1311 are by mills that are large. In constrast, only 467 of 3097 logger appearances are by large firms.

should submit higher bids, win disproportionately, and (likely) enter more often. To compare bids, we focus on the sealed bid auctions. We regress the per-unit bids (in logs) on a dummy for whether the bidder is a mill and auction fixed effects. The coefficient on the mill dummy is 0.239, meaning mill bids are 24% higher on average, with a t-statistic of roughly 7. An entering mill is also more likely to win than an entering logger (28% versus 21%).

To compare entry rates we require a measure of the potential number of mills and loggers that might enter a given auction. For this, we use the measure of "active" mills and loggers described above. The average sale had 5.1 potential mill entrants and 1.3 actual mill entrants, and 19.5 potential logger entrants and 3.0 actual logger entrants, so by this measure mills are more likely to enter.[21]

4. Comparing Auctions: Evidence

In this section, we investigate the consequences of auction choice for bidder participation, revenue and allocation. Our empirical approach is fairly straightforward; we describe it now before turning to the specific questions.

A. Empirical Approach

For a given outcome Y (such as the number of entering mills or loggers, or the auction price per unit), suppose that

$$Y = f(SEALED, X, \varepsilon), \tag{8}$$

where f is an unknown function, $SEALED$ is a dummy equal to one if the auction is sealed and zero if the auction is open, X is a vector of observed sale characteristics, and ε is unobservable. A standard point is that to identify the average effect of auction format, denoted $\tau_Y = \mathbb{E}_{X,\varepsilon}[f(1,X,\varepsilon) - f(0,X,\varepsilon)]$, we require that the unobserved component of the outcome is independent of the auction format conditional on covariates.

[21] Although (as noted above) our measure of active loggers is probably biased upward (relative to the measure for mills) due to difficulties in determining unique identities from abbreviated bidder names, it is unlikely the bias could be large enough to fully account for the difference.

This identification condition clearly holds for the randomly assigned sales in our sample (although it is important that the administrative unit that assigned the format is included in X, given that assignment probabilities differed by forest district).[22] It holds for the other sales if the forest manager's choice of format is based on information from the Forest Service appraisal, or follows some rule based on covariates in our data.[23]

Perhaps the most obvious approach to estimating τ_Y is to use ordinary least squares regression for the specification

$$Y = \alpha \cdot SEALED + X\beta + \varepsilon. \tag{9}$$

This approach is easily interpretable, but there are caveats. First, (9) does not allow the effect of sealed bidding to vary across tracts. To remedy this, we also report estimates from a specification where we interact $SEALED$ with the individual covariates.[24] A second issue is that we must specify the functional form for the covariates (and interactions among covariates) that will be included in X, but we have limited flexibility in doing so given our sample size relative to the number of covariates. While our results are not very sensitive to the alternatives we have tried, in general of course mis-specification could lead to bias.[25]

[22]Otherwise, we would be in danger of over-estimating the effect of sealed bidding if, for example, a forest district with especially valuable tracts also used a high fraction of sealed-bid sales. This is a shortcoming of Shuster and Nicolluci (1993)'s analysis: they control for only a limited set of tract characteristics, and so even for the randomized sales, the estimates they provide may not represent the causal effect of the auction format.

[23]If the forest manager uses a deterministic rule, such as using an open auction if and only if the volume of timber exceeds a threshold (which seems a possible description of some areas in California), then in principle auction format will not vary conditional on X. In practice, if our specification of X does not exactly match the rule, we will estimate $\Pr(SEALED|X)$ to be intermediate for sales close to the cut-off. So long as unobserved sale chacteristics are independent of the assignment conditional on X, we will still be identified in a manner analogous to a "regression discontinuity" approach, whereby discontinuous changes in the outcomes in response to changes in x close to the threshold will be attributed to auction format.

[24]We implement this approach by de-meaning the elements of X before interacting them with $SEALED$, so that the coefficient on $SEALED$ gives the average effect on the sample.

[25]There are really two concerns. First, if the covariates associated with open and sealed sales are fairly different, we will rely on our functional form assumptions to extrapolate what the outcome in one format would have been, had the auction been held using the other format. This concern motivates the procedure of selecting a subsample of sales with intermediate propensity scores. Second, if for instance sale volume is correlated with the auction format, a failure to flexibly control for sale

Motivated by this concern, we also report a set of estimates using a matching estimator. Because the matching estimator gives consistent estimates using a different approach than OLS, it provides a useful robustness check. This estimator matches every sealed bid auction with the M "closest" open auctions and vice versa, with closeness being measured as a weighted distance between sale characteristics.[26] It then compares the outcome of each sale t, Y_t, with the average outcome of the matched sales $\hat{Y}_t$, and estimates the average effect of auction format as the average of these comparisons:

$$\hat{\tau}_Y = \frac{1}{N^s} \sum_{t:sealed} (Y_t - \hat{Y}_t) + \frac{1}{N^o} \sum_{t:open} (\hat{Y}_t - Y_t),$$

where N^s and N^o are the number of sealed and open sales. We implement this estimator, setting $M = 4$, and compute robust standard errors following Abadie and Imbens (2004).

B. Evidence from Northern Forests

We begin our empirical analysis by looking at how auction choice affects the entry patterns of mills and loggers in the Northern forests. The model suggests that controlling for sale characteristics there should be more entry by loggers and either the same or less entry by mills. Table 3A reports our estimates (as well as our estimates of how auction choice effects other outcomes).

Conditional on sale characteristics, we estimate that sealed bid auctions attract 10-16% more logger entrants than open auctions. This translates roughly into 3-4 additional loggers for every 10 sales. All three point estimates are highly significant. In contrast, sale format appears to have little effect on entry by mills. Conditional on sale characteristics, our estimated effect is small and statistically cannot be distinguished from zero in all specifications.

The third column of Table 3 reports estimates of how auction format affects the fraction of entrants who are loggers. Consistent with the entry results, the composition of bidders at sealed bid auctions is shifted toward loggers. On average the

volume might lead us to falsely impute a revenue effect of auction method.

[26]We use the metric $||x||_W = (x'Wx)^{1/2}$, where W is a diagonal matrix consisting of the inverses of the variances of the covariates x. Thus the distance between two vectors of covariates x and z is $||x - z||_W$. We include the estimated propensity score for each auction as a covariate in addition to our standard set of characteristics.

fraction of participants who are loggers is 5-8% higher in sealed bid auctions than in open auctions.

Given this shift in bidder composition, it is natural to expect that sealed bid auctions should be more likely to be won by loggers. The fourth column of Table 3 reports our estimate of this effect. Our point estimates range from a 3.4%-7.4% greater chance that a logger will win if the auction is not sealed bid. These estimates are at best marginally statistically significant. Thus, although our point estimates are not insubstantial, we cannot rule out the effect of auction format on allocation being relatively small in these sales.

Finally, we turn to revenue, the issue that has attracted the most attention among economists. The fifth column of Table 3A reports our estimates of the effect of auction format on the sale price per unit volume. We find that after controlling for sale characteristics, sealed bid prices are 14-18% higher than open auction prices. Again, all three point estimates are all highly significant. To get a sense of the magnitude of this effect in dollar terms, note that the average winning bid (in 1983 dollars rather than 1983 dollars per unit volume) is just over $\$144,000$. So a 14% difference in the winning bid price translates into a $\$20,000$ difference in Forest Service revenue per sale, or about $19 million for the whole sample.

A natural question is whether the revenue difference is due to sealed bid auctions attracting more bidders. The final column of Table 3A reports estimates of the sale price where we include the number of entering loggers and mills as covariates. Even controlling for the number of entrants, sale method appears to matter. In the regression estimates, sealed bid auctions generate roughly 7% (s.e. 3%) more revenue. The matching estimator suggests a slightly larger revenue effect of 13% (s.e. 5%). The table does not report the revenue decomposition, but the estimates suggest that an additional mill is associated with about a 19% increase in the winning bid, while and additional logger is associated with about a 12% increase in the winning bid. Note that some caution is warranted in interpreting this revenue decomposition because there may be sale characteristics that are observed by the bidders but not accounted for in our data. In this case, the number of entrants may be endogenous in this regression.[27]

[27] An approach followed in the auction literature is to instrument for the number of entering bidders using measures of potential competition. We experimented with this, but found that our estimated

C. Evidence from California Forests

While the Northern forests seem particularly well-suited to making a statistical comparison between auction methods, we would like to draw on additional evidence as well. To this end, we also examined sales from California forests in the Forest Service's Pacific Southwest Region. We consider sales that took place between 1982 and 1989. We have data on 1188 open auctions and 694 sealed bid auctions.

While the Forest Service sale process is similar in California and the set of potential bidders includes both firms with manufacturing capability and logging companies, this sample is somewhat less ideal. The reason, which can be seen in the summary statistics in Table 1B, is that the tracts sold by sealed bid auction tend to be quite different from those sold by open auction. The principal difference is in the size of sales. The average sale volume for the open auctions is over 6000 mbf, while it is closer to 700 mbf for the sealed bid auctions. The sealed bid auctions are also more likely to be salvage sales. The per unit reserve prices are similar across sale formats.

The second column of Table 2 reports a logit estimate of the choice of sale method, using our standard controls. As is apparent in the summary numbers, volume is a highly important correlate of sale method. Sale method also varies significantly across the twelve forests in the region. The extent to which sale method correlates with sale characteristics can also be seen in Figure 1B, where we plot the density of the propensity score for the open and sealed bid auctions. Our logit regression predicts the sale method of many of the open auctions with near-perfect precision; this is mainly a function of the fact that very large sales are almost certain not to be sealed bid.

As with the Northern forests, we again drop sales that have an estimated propensity score below 0.075 and above 0.925. This dramatically reduces the sample and leaves us with 212 open auctions and 269 sealed bid auctions. Figure 1B illustrates how, relative to the full sample of California sales, the selected sample has much more overlap in the distribution of estimated propensity scores. And as can be seen in Table 1B, the selected sample has much smaller differences across sale format. Still, the remaining differences require carefully controlling for covariates in estimating the effect of auction format on different outcomes.

---

coefficients were highly sensitive to the particular choice of potential competition measures, none of which are ideal.

With this caveat in mind, we turn to Table 3B, where we report estimates of the effect of auction method on entry, revenue and allocation outcomes. The results for entry are similar to the Northern forests. Sealed bid auctions attract more loggers. The regression models give an estimate of 11-12% (6%) more loggers at sealed sales, which translates into an additional 3 loggers participating for every 10 sales (similar to the Northern forests in terms of the absolute numbers). The matching estimate is a bit larger — 4.7 additional loggers for every 10 sales. We also find that mills are somewhat less likely to participate in sealed bid sales. Our point estimate from the regression model is that sealed bidding attract 1.3 fewer mills for every 10 sales, but the estimate is not statistically significant. The matching estimate is larger in magnitude: 3 fewer mills for every 10 sales, and this estimate is statistically significant. As in the Northern forests, the composition of bidders shifts significantly toward logging companies with sealed bidding — here by 8-15%.

Our estimates of the effect of auction method on allocation also are qualitatively similar those in the Northern forests, though larger and more significant. In the California forests, we estimate that there is roughly a 8-14% greater chance a logger will win with sealed bidding (the linear probability estimate is 8% (5%), the matching estimate is 14% (5%)).

A notable difference between the California results and those for the Northern forests is that we do not find a significant effect of auction method on revenue in California. The regression estimate is slightly positive, the matching estimate slightly negative. Neither are large or statistically insignificant, and the same is true after controlling for the number of entering mills and loggers.

D. Explaining the Departures from Revenue Equivalence

Our empirical evidence suggests that in both the Northern and California forests there are significant differences between the outcomes of sealed bid and open auctions. Conditional on sale characteristics, sealed bid auctions attract more entry by logging companies, with either a negligible change in the entry of mills (Northern region) or a decrease in their participation (California). Sealed bidding also appears more likely to result in the auction being won by a logging company — particularly in California. Finally, after controlling for sale characteristics, the winning bids in the sealed bid sales are appreciably higher in the Northern forests (14-17%), but similar to open auction prices in California. It is in the effect of auction method on sale price that the two regions differ most noticeably.

At a qualitative level, the theoretical model developed earlier in the paper can rationalize all of these findings. The model predicts that logger entry will be higher in sealed bid sales, that loggers are more likely to win a sealed bid sale, and that sealed bid sales may result in greater revenue, particularly if mills are able to collude to some extent in the open sales. Moreover, the key assumption generating these departures from revenue equivalence, that bidders are heterogeneous, also seems consistent with the data.[28]

What we cannot say at this point, however, is whether a reasonable parametrization of the model can match our quantitative findings. Moreover, recall that the theory predicts qualitatively the same differences between open and sealed bidding regardless of whether the mills are able to collude in open auctions, a primary concern that has historically motivated the use of sealed bidding in Forest Service timber auctions. Without a more quantitative approach to the model, we cannot distinguish between its competitive and collusive versions. We try to address this shortcoming in the next section by estimating the model's parameters directly from the data and then comparing the quantitative predictions of the theories to the data.

E. Alternative Explanations

A different explanation for our findings is that our estimates do not reflect the systematic effects of auction format, but rather some confounding correlation between auction choice and unobserved aspects of the sale that also affect the outcome. This is certainly a concern. Even in the Northern forests, where many sale assignments were random, we may not have perfectly controlled for sale differences. And as we have noted the differences are greater in California. We have attempted to mitigate this by making use of the very rich data on sale characteristics in the Forest Service sale reports, augmented by further data on market conditions.

Could it be the case that some omitted variable is generating our findings? Several of the most obvious stories have problems themselves. For instance, one possibility

[28] Above, we reported comparisons between mills and loggers for the Idaho and Montana sales. In California, mill bids are just over 10% higher on average, after controlling for auction fixed effects, and the difference is highly significant. Mills are also more likely more likely to participate and to win conditional on participating.

is that forest managers like to sell more valuable tracts by sealed bid, a bias that would help to explain the entry and revenue differences we find. This story is hard to square, however, with the fact that larger sales, which are by definition more valuable on a total value basis, are more often sold by open auction. A second possibility is that forest managers use sealed bid sales when they expect there to be more bidder interest, especially on the part of logging companies. This would certainly help to explain the entry results, though it is not clear to us why forest managers would systematically behave in this way. Indeed, industry lore is more consistent with a scenario where the mills prefer oral auctions (as predicted by our theory), and where forest managers defer to the mill's preferences.[29]

Turning from endogeneity to behavioral explanations, recall that our theoretical model abstracted from two potentially relevant aspects of timber auctions: common values and bidder risk-aversion. Could either of these explain our empirical findings? While our results certainly do not rule out the presence of common values or bidder risk-aversion (or both), it seems unlikely that either is primary source of the departures we observe from revenue equivalence. With common values (and without the other elements of our model, namely bidder heterogeneity and collusion), prices should be *lower* in sealed auctions, rather than higher as we observe in the data. Risk-aversion might be able to explain the observed prices, but it would also suggest that participation should be *lower* in the sealed bid auctions, rather than higher. So to the extent that either common values or bidder risk-aversion would help to explain our findings, they would have to be part of a more complicated story.

6. Structural Estimation and Testing

In this section, we try to assess more precisely the relationship between our findings and the theory we proposed to account for them. We investigate three related issues. First, we ask whether our a calibrated version of our model, with parameters estimated from the data, can quantitatively match the departures we observe from revenue equivalence. Second, we ask whether the model can provide a measure of

[29] A further point that is potentially relevant here that the Forest Service Handbook instructs forest managers to use sealed bidding if they expect a sale *not* to be competitive (CITE). To the extent that this might create an omitted variable problem, it presumably would lead to finding that sealed bidding generated *less entry* and *lower prices*, precisely the opposite of our results.

bidder competitiveness in the open auctions. Finally, we estimate the welfare consequences of moving exclusively to open or sealed bidding, under the assumption that our estimated model accurately describes the sale environment.

We proceed as follows. We first use the sealed bid data to estimate the unknown parameters of the theoretical model — the value distributions of loggers and mills, and the costs of entry. We then compute the expected equilibrium outcome of each sale as predicted by our calibrated model and compare the predicted outcomes to what we actually observe. This allows us to assess how accurately the model predicts the observed differences across auction formats, and, as we will see, provides a rough way to assess bidder competitiveness in the open auctions. The welfare comparison of open and sealed bidding is developed at the end of the section.

A. Structural Estimation

Our first step is to use the sealed bid data to estimate the parameters of the theoretical model. To estimate the value distributions of mills and loggers, we follow the two-stage approach pioneered by Guerre, Perrigne and Vuong (2000). They suggest fitting a distribution to the observed sealed bids, then using the first-order condition for optimal bidding to recover the bidders' value distributions. This approach parallels the recovery of latent cost parameters from price-quantity data (e.g. Rosse, 1976). Given the value distributions, we can proceed to estimate bidders' entry costs.

We begin by introducing some notation. For a given auction, let X denote the set of sale characteristics known both the us and the bidders. To account for the fact that sales may differ in ways we cannot observe, we let u denote an additional characteristic known to participating bidders but not observed in our data. We write the bidders' value distributions, conditional on (X,u), as $F_L(\cdot|X,u)$ and $F_M(\cdot|X,u)$.

In line with our model, we assume that bidders' values, and hence their bids, are independent conditional on (X,u). Letting n denote the set of participating bidders, we write the equilibrium bid distributions as $G_L(\cdot|X,u,n)$ and $G_M(\cdot|X,u,n)$. We assume that if there is a single bidder, he optimally bids the reserve price, but otherwise treat the reserve price as non-binding.[30] Finally, we assume that bidders

[30]See Haile (2001) for a discussion of why the reserve prices in Forest Service timber auctions are typically non-binding. A slight drawback to our treatment of the reserve price is that our fitted bid distributions will assign positive (though typically small) probability to bids below the reserve price.

have independent entry costs that are concentrated around some average entry cost $K(X)$.

Estimating the Bid Distributions

Conditional on the observable sale characteristics (X,n), the joint distribution of bids in a given auction is a combination of the bid distributions $G_L(\cdot|X,u,n)$ and $G_M(\cdot|X,u,n)$ and the distribution of the unobserved auction heterogeneity u, which is responsible for any covariation of bids within a given auction. We adopt a parametric approach to estimate these three distributions.

After extensive experimentation, we found that Weibull bid distributions with Gamma distributed auction heterogeneity provided a good fit to our sealed bid data. Thus we assume that for $k = L,M$:

$$G_k\left(b|X,u,n\right) = 1 - \exp\left(-u \cdot \left(\frac{b}{\lambda_k(X,n)}\right)^{p_k(n)}\right).$$

Here $\lambda_k(\cdot)$ is the scale, and $p_k(\cdot)$ the shape, of the Weibull distribution, parametrized as $\ln \lambda_k(X,n) = X\beta_x + n\beta_n + \beta_k$ and $\ln p_k(n) = n\gamma_n + \gamma_k$.[31] We assume u has a Gamma distribution with unit mean and variance θ, and is independent of X and n.[32] We estimate the parameters $(\beta_x, \beta_n, \beta_L, \beta_M, \gamma_n, \gamma_L, \gamma_M, \theta)$ by maximum likelihood; the estimates are reported in Table 4.

Several points about the estimated bid distributions deserve mention. First, recall that the basic assumption of the theory was that mill values stochastically dominate logger values, and an implication was that mill bids should dominate logger bids. Our empirical specification does not impose this, nor any ordering between logger and mill bids. Nonetheless, we find that mill bids do dominate those of loggers. On average, mill bids are roughly 25% higher than logger bids. Also consistent with the theoretical model, we find that bids are increasing in the number of competitors

[31] The specification for X we adopt is more parsimonious than in our earlier regressions. While our results do not seem sensitive to including more covariates (or less for that matter), we opted for parsimony because of the need to make out-of-sample predictions where over-fitting could in principle be a problem.

[32] Implicitly then, u is observed only once bidders acquire information. The assumption that u is orthogonal to (X,n) is strong, but should be viewed in light of most empirical work on auctions (an important exception being Krasnokutskaya, 2003), which makes the even stronger assumption that there is no unobserved heterogeneity at all across auctions.

(a property that may not be satisfied by auction models where bidder values are affiliated or have a common value component). Finally, we estimate that u has significant variance, indicating that our modeling of the unobserved heterogeneity across auctions is warranted.

Estimating the Value Distributions

Given our fitted bid distributions, we turn to recovering the bidders' value distributions. Under the assumption that the observed bidding is consistent with equilibrium behavior, each bid must be optimal against the opponents' bid distributions. So given sale characteristics (X,u,n), a bidder's value v_i can be recovered from his observed bid b_i and his first-order condition for optimal bidding:

$$\frac{1}{v_i - b_i} = \sum_{j \in n \setminus i} \frac{g_j(b_i; X,u,n)}{G_j(b_i; X,u,n)}. \tag{10}$$

Note that for a given auction in our data, we do not know u, so we cannot recover the value that corresponds to each observed bid. Nevertheless, as established by Krasnokutskaya (2002), we can recover the distributions $F_L(\cdot|X,u)$ and $F_M(\cdot|X,u)$ by sampling bids from $G_L(\cdot|X,u,n)$ and $G_M(\cdot|X,u,n)$ and inferring the corresponding values from the first-order conditions.[33,34]

[33]Guerre, Perrigne and Vuong (2000) and Krasnokutskaya (2002) suggest a nonparametric approach for both the first step of this approach (estimating the bid distributions) and the second step (estimating the distribution of valuations). In contrast (and similar to an approach used by Jofre-Benet and Pesendorfer (2003)), we use a parametric model for the first step, but use a nonparametric approach for the second step. We follow this approach because our sample size is small relative to the number of covariates and variation in participation.

[34]In recovering values from the first-order condition, we face a self-imposed technical hurdle. It can be shown that in a (symmetric) IPV auction model, $\mathbb{E}[v_i] = (\bar{b} + (n-2)\mathbb{E}[b_i])/(n-1)$. Because we model the bid distribution as Weibull, and the Weibull distribution has infinite support, we implicitly impose an infinite mean on the bidder value distribution. An obvious alternative would be to estimate bid distributions with finite support, but this has serious drawbacks as well because it requires estimating the maximum bid conditional on (observed and unobserved) covariates. This is a hard problem and moreover the mean of bidder values will be in close correspondence with the (arguably poor) estimate. Instead, we propose to exploit the fact the for high values, the bid function is very flat, so by specifying even a wide range for the maximum *value*, we can pin down the maximum bid with some precision. We have taken this approach in generating Figure 2. Tables 5 and 6, however, still reflect our initial approach, similar in spirit, which was to cap values in the far right tail at six times the bid draw (six is arbitrary, but using two or ten has little effect on our results). From what we have seen so far, our results do not seem very sensitive to the way that we impose the upper bound.

Figure 2 plots the density functions for logger and mill values for an auction with average covariates, as well as the equilibrium bid functions assuming two mills and two loggers participate in the auction. As the Figure indicates, the distribution of mill values is substantially shifted rightward from the distribution of logger values. Moreover, the estimated mill bid function is below the logger bid function. Thus, mills bid less than loggers for any given value, matching a key prediction of the theoretical model.

Estimating Entry Costs

The remaining parameter of the model is the entry cost. To recover the entry cost for each auction, we use the equilibrium entry condition. Recall that when entry costs are concentrated as we have assumed, and loggers are the marginal participants, then in equilibrium each logger will be close to indifferent regarding participation. So if the expected logger profit is $\Pi_L(X)$, the average entry cost can be recovered from the equilibrium condition: $K(X) \approx \Pi_L(X)$.

We write the expected profit for a participating logger, conditional on observed sale characteristics, as

$$\Pi_L(X) = \sum_{n \subset N} \Pi_L(X,n) \Pr[n|X,i \in n]. \tag{11}$$

The first term, $\Pi_L(X,n)$, is a logger's expected profit conditional on sale characteristics and the set of participants. This number is easily computed from our estimate of the value and bid distributions, integrating out the unobserved auction heterogeneity.

The second term, $\Pr[n|X,i \in n]$, is the probability an entering logger assigns to the set of participants being n. To estimate $\Pr[n|X,i \in n]$, we assume that bidders know the number of potential mill entrants (and hence the actual number because mills will be inframarginal participants in equilibrium). We model the number of entering loggers as a Poisson random variable with mean $\mu(X)$, parametrized as $\ln \mu(X) = X\alpha$.[35] Our estimate of $\mu(X)$ is reported in Table 4.

The final step, having estimated $\Pi_L(X,n)$ and $\Pr[n|X,i \in n]$, is to calculate for each sale the expected profit for a logger that enters a sealed bid auction for that

[35] In theory, the distribution of logger entrants is binomial because loggers make independent entry decisions. As we do not have a very good measure of the number of potential logger entrants, we use the poisson specification to approximate the binomial.

tract.[36] We then impute the entry cost for each tract t as $K(X_t) = \Pi_L(X_t)$. In our sample of sealed bid auctions, we estimate that the median entry cost is $4449 (s.e. $503). As the costs of surveying a tract can run to several thousand dollars, this seems reasonably consistent with our prior beliefs about the costs of acquiring information and preparing a bid.[37]

B. Comparing Predicted and Actual Outcomes

Having estimated the parameters of the theoretical model — the auction entry costs and the distribution of logger and mill values — we now ask how closely the model's equilibrium predictions match the observed auction outcomes in our data. This serves several objectives. In the case of sealed bid sales, this exercise provides a measure of how well we have fit the entry and bidding data. In the case of open auctions, it allows us to ask whether the calibrated model can explain (out-of-sample) the open auction outcomes, and in particular, whether assuming some degree of co-operative behavior provides a more accurate fit to the data. Finally, by looking at both kinds of sales, we can assess whether the model is able to explain not just the qualitative but the quantitative departures from revenue equivalence documented above.

To proceed, we compute for each tract the expected equilibrium outcome of our parameterized model. If the tract was sold by sealed bid auction, we compute the expected equilibrium outcome for a sealed bid auction. If the tract was sold by open auction, we compute both the competitive and collusive open auction equilibria. We report two variations of these calculations. The first is the expected equilibrium outcome of the model conditional on the actual participation in each sale. The second is the expected outcome of the full entry equilibrium.[38]

[36]In practice, of course, not every tract was sold by sealed bidding. Nonetheless, our estimates of the bid and value distributions, and logger entry under sealed bidding, combined with our knowledge of each tract's characteristics, allow us to calculate the expected logger profits of a hypothetical sealed bid auction for each tract, even though some tracts were actually sold by open auction.

[37]An alternative approach to estimating entry costs is to assume that bidders *perfectly* forecast opponent entry. Then in equilibrium, $\Pi_L(X,n) \approx K(X)$, where n is the *realized* entry. This approach generates a slightly lower estimate for the median entry cost in our sample of sealed bid auctions: $2652 (s.e. $334), but one that also seems within reason.

[38]The actual amount of "computation" is limited. In the case of the sealed bid auctions, the equilibrium bid and entry distributions correspond to our estimates from the previous section. In the open auctions, bidders either bid their value, or if they are colluding, potentially drop out

Table 6 reports the actual average outcomes in our sample and the average outcomes predicted by our parameterized model.[39] In the case of sealed bid auctions, we closely predict the average bids of loggers and mills. The model predicts averages of 54.1 (\$/mbf) and 102.4 (\$/mbf), while the actual averages are \$57.6 and \$101.0. It also closely predicts logger entry, as well as the average auction prices both unconditionally (i.e. given just sale characteristics) and conditional on the set of participating bidders. The model somewhat under-predicts the fraction of sales that loggers win — both the unconditional prediction of 63.5% and the prediction of 66.0% conditional on realized entry undershoot the actual number of 68.7%.

Of course, it should not be too surprising that the model accurately predicts the sealed bid outcomes because the its parameters are estimated from the sealed bid data. The more demanding test of how well the theory can fit the observed outcomes is to compare the open auction outcomes predicted by the model to the actual outcomes. In this case, we asking the model to make predictions that are "out-of-sample" in two senses: we are predicting sale outcomes for tracts not used to estimate the model's parameters, and also for a different auction format than that used to estimate the model's parameters.

To consider the open auctions, we start by looking at the model's predictions for entry and allocation relative to the realized outcomes. Strikingly, the model predicts a level of equilibrium logger entry that is very close to the level we actually observe (2.81 loggers per sale versus 2.84 in reality), indicating that the fitted model is able to explain the entry differences between open and sealed bid sales in our data. The model is somewhat less successful in matching the fraction of sales won by loggers. As with the sealed bid auctions, the model under-predicts how often loggers win (the model's prediction is that loggers will win 51.5% of the sales, or 54.5% conditional on realized participation, while in reality they win 60%). Note, however, that despite under-predicting logger purchases for both sale formats, the model accurately captures the

immediately. Moreover, with concentrated entry costs, all potential mills enter, meaning the only unknown is the entry behavior of loggers at open auctions. To compute this, we approximate their equilibrium entry distribution as Poisson and for each open auction solve for the Poisson parameter that just equates expected logger profits and our estimate of the entry cost. Logger entry is identical under competition and collusion with concentrated entry costs, so one computation per auction suffices.

[39]We generate the standard errors using a parametric bootstrap in which we re-sample from the asymptotic distribution of the bid and entry distribution parameters reported in Table 4.

difference across the open and sealed bid sales (the model's prediction is 8.7% versus 12% in reality).

A key point to have in mind for the open auctions is that under our assumption of concentrated entry costs, the competitive and collusive equilibria differ *only* in the price they predict. Therefore, to the extent that the data squares with the model in terms of entry and allocation, it squares with many versions of the model — from competitive behavior to perfect mill collusion, to any intermediate degree of mill collusion that involves the highest-valued mill bidding up to his value. To distinguish these behavioral alternatives and extract some measure of competitiveness, it is necessary to focus on prices.

The numbers in Table 6 indicate that the observed prices in the open auctions lie between the competitive and fully collusive prices predicted by the model. In reality, the average sale price across open auctions in our sample is \$72.8 per mbf; the competitive model predicts average prices equal to \$79.0, or \$79.8 if we condition on realized entry. In contrast, if mills fully collude, we predict the average price will be \$51.5 per mbf. Even accounting for sampling error, we reject both the competitive and collusive models at the 1% level. Thus it appears that the assumption of mildly cooperative behavior on the part of participating mills provides a better match than either the competitive or fully collusive extremes.

When we put the sealed bid and open auction comparisons together, it appears that the theoretical model developed in Section 2 and estimated using the sealed bid data does a plausible job of explaining the departures from revenue equivalence we observe in the data. The parametrized model easily explains the logger and mill entry patterns we observe in the Northern sales. It errs somewhat in predicting the fraction of sales won by loggers for both the sealed bid and the open sales, but predicts a difference across auction formats that is similar to what we actually observe. Finally, the observed price differential between the open and sealed bid sales in our sample allows us to reject (under the assumption that sealed bidding behavior is competitive) both the assumption of perfectly competitive behavior and perfectly collusive mill behavior in the open auctions. Rather, the price difference can be rationalized by a mild degree of cooperation on the part of mills bidding in open auctions.

C. Quantifying the Trade-offs in Auction Design

So far our goal has been to assess whether the theoretical model we proposed could explain the systematic departures from revenue equivalence we observe in the data. We now take as given that we have accurately estimated bidders' values and entry costs, and investigate the welfare consequences of using either open or sealed bidding on an exclusive basis. From an a priori standpoint, our theoretical results suggest that neither format will dominate. The open auction conveys an efficiency benefit, but the increase in social surplus may come at the cost of lost revenue and an allocation that favors stronger bidders. For this reason, it seems natural to try to quantify the trade-offs faced in choosing between the two formats.

To estimate the welfare consequences of making exclusive use of either open of sealed bidding, we compute the predicted outcome of both an open and a sealed bid auction for each tract in our sample. We consider two alternative specifications of mill behavior in the open auctions: a benchmark specification where mills behave competitively, and perhaps a more realistic specification where they cooperate 25% of the time (25% being the number that rationalizes the observed open auction prices).

Our comparisons reported in Table 7. The top panel reports the expected auction outcomes taking participation as fixed and computing only the corresponding bidding equilibrium. The bottom panel reports expected outcomes when we solve for the complete entry equilibria of the alternative models.

A first point that stands out is that if participation is assumed to be independent of the auction format, the differences in equilibrium outcomes between open and sealed bidding — assuming bidder behavior is competitive in both cases— are very small. Sealed bidding would generate more revenue, but the revenue gain is less than $1000 per sale. Sealed bidding also increases the probability that sales are won by loggers, but again the average increase in probability is less than 1%. Finally, the efficiency benefit to using an open auction format is also quite small, only $145 per sale.

These differences are magnified when we account for the fact that bidder participation will vary systematically with auction format. Under our assumption of concentrated entry costs, sealed bid and open auctions will attract the same number of mills, but sealed bid auctions will attract between 3-4 more loggers for every 10 sales. This additional entry has several effects. First, it generates a more substantial difference in the probability that a logger will win any given sale — the probability a given sale will be won by a logger is 3.2% higher with sealed bidding. Second, it increases the revenue advantage of sealed bidding. With equilibrium participation, we estimate that sealed bidding would raise an extra $8000 on average (about 5% of the average sale price). At the same time, however, the additional entry has a welfare cost: on a per-sale basis, the social surplus generated by open auctions is nearly $1900 higher than with sealed bidding.

As a practical matter, however, the model suggests that these differences are dwarfed by the potential effects of bidder collusion. Even if we take participation as fixed, open bidding generates some $22,000 less per sale than competitive sealed bidding if mills are able to engage in a mild amount of cooperative behavior. The difference is over $29,000 once we account for participation effects. So to the extent that mild cooperation by mills at open auctions is the behavioral assumption that receives the most support from our data, the revenue benefits of sealed bidding clearly seem to be the most quantitatively significant welfare consequence of the choice of auction method — at least in these forests during the study period.

A more general point we draw from this is that the degree of bidder competitiveness can be of crucial importance in weighing the relative benefits of open and sealed bidding. Indeed, we have repeated the analysis of this section on the California forests with quite different conclusions. Recall that in the those forests, auction format appeared to have an insignificant effect on prices. Though the greater heterogeneity between open and sealed sales makes the California analysis somewhat less precise, when we calibrate the model using the sealed bid sales, we find that the observed auction open behavior can be described reasonably well under the assumption that firms bid competitively. As a consequence, the welfare analysis for California resembles the comparison of competitive sealed and open auctions in Table 7, with moderate efficiency gains for open sales trading off against moderate revenue benefits to sealed bidding.[40]

6. Conclusion

This paper has examined the relative performance of open and sealed bid auctions, using U.S. Forest Service timber sales as a test case in auction design. Our main em-

[40]We view the California analysis as still having some room for improvement; we may report more fully on it in the next version of the paper.

pirical finding is that there are systematic differences in the auction outcomes. Sealed bid auctions attract more small bidders, shift the allocation toward these bidder, and in the Northern forests though not in California forests, generate higher revenue. Our main theoretical contribution is an extension of the standard independent private values auction model that can explain these findings, both qualitatively and quantitatively, and also allows us to measure the degree of bidder competitiveness. Our structural estimation results suggest that competitiveness may vary across Forest Service regions, and that the implications of competitiveness for auction choice may be quantitatively the most significant.

Appendix: Proofs of the Results

To begin, we establish existence of entry equilibrium.

Proposition 7 *For both auction formats, a type-symmetric entry equilibrium exists.*

Proof. For the sealed bid auction, Li and Riley (1999) show that for any set of participants, there is a unique type-symmetric bidding equilibrium that is type-symmetric. The same is true for the open auction if we restrict attention to undominated strategies. We can use a single proof to show the existence of an entry equilibrium for both auction formats. Let $\Pi_i(K)$ denote i's profits from entry assuming entrants use equilibrium bid strategies. An entry equilibrium couples these strategies with a vector K such that $\pi_i(K) := \min\{\Pi_i(K), \overline{k}\} = K_i$ for all i. So establishing a type-symmetric entry equilibrium amounts to finding a type-symmetric fixed point of $\pi = (\pi_1,...,\pi_n)$. Let $\mathcal{K} = \{(K \in [\underline{k}, \overline{k}]^n : K_1 = ... = K_L, K_{L+1} = ... = K_N\}$ denote the space of type-symmetric entry thresholds. Now, $\pi : \mathcal{K} \to \mathcal{K}$ and is continuous in K because $\Pi(\cdot)$ results from a unique Nash equilibrium (and hence is continuous in K). So Kakutani's fixed point theorem implies that π has a fixed point in $\mathcal{K}$. *Q.E.D.*

Proof of Proposition 1. Equilibrium existence is shown above. Properties *(i)* and *(ii)* follow from the analysis of Maskin and Riley.

Proof of Proposition 2. Equilibrium existence is shown above. Properties *(i)* and *(ii)* follow from the fact that it is a dominant strategy for participants to bid their values, and by Assumption *(ii)*, $F_M(b) \leq F_L(b)$ for all b.

Proof of Proposition 3. Standard revenue equivalence results (see e.g. Milgrom, 2003) imply that for any fixed set of participants n, the equilibrium surplus, and the expected revenue and profits of individual bidders will be identical across the two auction formats. Therefore $\Pi_i^s(K) = \Pi_i^o(K)$. Moreover, $\Pi_i = \Pi_i^o = \Pi_i^s$ is constant in K_i and decreasing in K_j, so it is decreasing in K when $K = (K,...,K)$. It follows that both auctions have a unique symmetric entry equilibrium; the equilibrium entry threshold solves $\Pi_i(K,...,K) = \min\{\overline{k},K\}$. The results follow directly.

Proof of Proposition 4. We first show that for any K,

$$\Pi_L^s(K) \geq \Pi_L^o(K) \qquad \text{and} \qquad \Pi_M^s(K) \geq \Pi_M^o(K).$$

From the text, for each auction type $\tau \in \{s,o\}$:

$$\Pi_i^\tau(v_i;n) = \int_{\underline{v}}^{v_i} \Pr[i \text{ wins} \mid v_i = x,n,\tau\,]dx.$$

In the sealed bid equilibrium $b_M(v;n) \leq b_L(v;n)$ for all v, while all bidders use the same strategy in the open auction. Therefore if i is a logger:

$$\Pr[i \text{ wins} \mid v_i,\ n,\ \text{Open}] \geq \Pr[i \text{ wins} \mid v_i,\ n, \text{Sealed}],$$

while the converse holds for mills. Hence $\Pi_L^s(v;n) \geq \Pi_L^o(v;n)$ and consequently $\Pi_L^s(K) \geq \Pi_L^o(K)$, while the converse holds for mills.

To proceed, we characterize type-symmetric entry equilibria of the open auction. Define:

$$\pi_i(K) := \min\{\Pi_i^o(K), \overline{k}\}.$$

Consider the space $\{(K_L,K_M) : K_i \in [\underline{k},\overline{k}]\}$ of type-symmetric entry thresholds. Let $\mathcal{L}_L$ denote the locus of points (K_L,K_M) for which $\pi_L(K) - K_L = 0$, and define $\mathcal{L}_M$ accordingly. The intersections of $\mathcal{L}_M$ and $\mathcal{L}_L$ are the type-symmetric entry equilibria of the open auction. Figure 1 depicts a unique equilibrium, but there may be several.

Now, observe that $\mathcal{L}_L$ and $\mathcal{L}_M$ are continuous, and also downward sloping because $\Pi_i^o(\cdot)$ and hence $\pi_i(\cdot)$ are decreasing in K. So above $\mathcal{L}_i$, $\pi_i(K) - K_i < 0$ while below $\mathcal{L}_i$, $\pi_i(K) - K_i > 0$. Moreover, if $K_M = \overline{k}$ then by definition $\pi_i(K) - K_i \leq 0$, so $\mathcal{L}_M$ lies (weakly) below $\mathcal{L}_L$ at $K_L = \overline{k}$. A consequence is that for any point $K = (K_M,K_L)$ above $\mathcal{L}_L$ and below $\mathcal{L}_M$, there must be an open auction equilibrium K^o with $K_M^o \geq K_M$ and $K_L^o \leq K_L$.

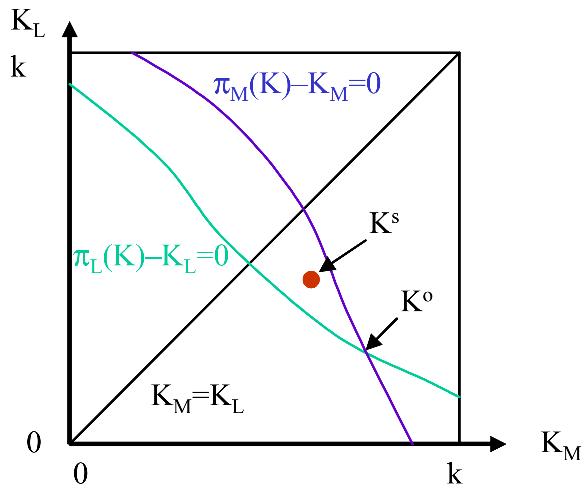

Graphical Depiction of Entry Equilibrium

To prove the result, suppose K^s is a type-symmetric entry equilibrium of the sealed auction and is interior (a similar argument applies for boundary equilibria). Because mills prefer open auctions and loggers sealed auctions:

$$\pi_M(K^s) - K_M^s \geq 0 \qquad \text{and} \qquad \pi_L(K^s) - K_L^s \leq 0.$$

So K^s lies above $\mathcal{L}_M$, below $\mathcal{L}_L$ (as in the Figure). Therefore there is a type-symmetric open auction equilibrium K^o with $K_M^o \geq K_M^s$ and $K_L^o \leq K_L^s$. Relative to the sealed equilibrium, mills enter more, while loggers enter, and win, less. *Q.E.D.*

Proof of Proposition 5. Social efficiency means choosing entry thresholds and an allocation process to maximize social surplus. Given participation, it is best to allocate efficiently, which the open auction does and the sealed bid auction already fails to do. To consider entry, write the social surplus as $S(K_1,...,K_N)$. A participant's open auction payoff equals his contribution to social surplus, so bidder payoffs in the open auction entry game can be written:

$$u_i(K_1,...,K_N) := \Pi_i(K_1,...,K_N) - K_i = S(K_1,...,K_N) - h_i(K_{-i}),$$

and satisfy $du_i/dK_i = dS/dK_i$.

Now, suppose K^{**} is an efficient type-symmetric profile:

$$K^{**} = \arg\max_{K_L,K_M} s(K_L,K_M) = S(K_L,...,K_L,K_M,...,K_M).$$

Assuming S is interior (the proof is similar for boundaries), we have $ds/dK_L(K^{**}) = ds/dK_M(K^{**}) = 0$. But then, because S is symmetric in logger thresholds and in mill thresholds, it must be that for all i, $dS/dK_i(K^{**}) = 0$ and hence $du_i/dK_i(K^{**}) = 0$. So K^{**} is also a type-symmetric entry equilibrium. *Q.E.D.*

Proof of Proposition 6. Let $\Pi_i^c(K)$ denote the profits of bidder i from entering if mills collude. We have:

$$\Pi_L^c(K) = \Pi_L^o(K) \qquad \text{and} \qquad \Pi_M^c(K) \geq \Pi_M^o(K).$$

Now consider the depiction of equilibrium open auction entry in the Figure above. Collusion by mills has the effect of increasing mill profits for any (K_L, K_M) pair, so the curve $\mathcal{L}_M$ shifts up, while $\mathcal{L}_L$ stays unchanged. Because $\mathcal{L}_M$ must still lie below $\mathcal{L}_L$ when $K_M = \overline{k}$, this means that for any open auction entry equilibrium, there must clearly be a collusive equilibrium with more mill entry, less logger entry and less chance of a logger winning. *Q.E.D.*

References

Abadie, A. and G. Imbens, "Large Sample Properties of Matching Estimators for Average Treatment Effects," UC Berkeley Working Paper, 2004.

Athey, S. and J. Levin, "Information and Competition in U.S. Forest Service Timber Auctions," *Journal of Political Economy*, April 2001, 109(2), pp. 375–417.

Bajari, P., "The First-Price Auction with Asymmetric Bidders: Theory and Applications," Ph.D. Dissertation (University of Minnesota), 1997.

Bajari, P. and L. Ye, "Deciding Between Competition and Collusion," *The Review of Economics and Statistics*, November 2003, pp. 971-989.

Bajari, P. and G. Summers, "Detecting Collusion in Procurement Auctions," *Antitrust Law Journal* 70(1), 2002, 143-170.

Baldwin, L., R. Marshall, and J.-F. Richard., "Bidder Collusion at U.S. Forest Service Timber Auctions," *Journal of Political Economy*, August 1997, 105(4), pp. 657–699.

Brannman, L., and L. Froeb, "Mergers, Cartels, Set-Asides and Bidding Preferences in Asymmetric Second-price Auctions," *Review of Economics and Statistics*, 82(2) (2000) 283-290

Brendstrup, B. and H. Paarsch, "Nonparametric Estimation of Dutch and First-Price, Sealed-Bid Auction Models with Asymmetric Bidders." Working Paper, University of Iowa Department of Economics, 2003a.

Brendstrup, B. and H. Paarsch, "Nonparametric Identification and Estimation of Multi-Unit, Sequential, Oral, Ascending-Price Auctions with Asymmetric Bidders." Working Paper, University of Iowa Department of Economics, 2003b.

Guerre, E., I. Perrigne and Q. Vuong, "Optimal Nonparametric Estimation of First Price Auctions," *Econometrica*, May 2000, 68(3), pp. 525–574.

Froeb, L. and P. McAfee, "Deterring Bid Rigging in Forest Service Timber Auctions," U.S. Department of Justice Economic Analysis Group Working Paper, 1988.

Haile, P., "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales," *American Economic Review*, June 2001, 92(3), pp. 399–427.

Haile, P., H. Hong and M. Shum, "Nonparametric Tests for Common Values in First-Price Sealed Bid Auctions," Yale University Working Paper, 2003.

Hansen, R., "Sealed-Bid Versus Open Auctions: The Evidence," *Economic Inquiry*, January 1986, 24(1), pp. 125–142.

Jofre-Bonet, M. and M. Pesendorfer, "Estimation of a Dynamic Auction Game," Econometrica, 2003, 71: 1443–1489.

Johnson, R., "Oral Auctions versus Sealed Bids: An Empirical Investigation," *Natural Resources Journal*, April 1979, 19(2), pp. 315–335.

Krasnokutskaya, E., "Identification and Estimation in Highway Procurement Auctions with Unobserved Heterogeneity," Yale Working Paper, 2002.

Li, H. and J. Riley, "Auction Choice," UCLA Working Paper, 1999.

Matthews, S. "Comparing Auctions for Risk Averse Buyers: A Buyer's Point of View," *Econometrica*, May 1987, 55(3), pp. 633-646.

Maskin, E. and J. Riley, "Asymmetric Auctions," *Review of Economic Studies*, July 2000, 67(3), pp. 413–438.

Mead, W., "Natural Resource Disposal – Oral Auction Versus Sealed Bids," *Natural Resources Journal*, April 1967, 7(2), pp. 194–224.

Milgrom, P., *Putting Auction Theory to Work*, Cambridge University Press, 2004.

Perrigne, I., "Random Reserve Prices and Risk Aversion in Timber Sale Auctions," Mimeo, Pennsylvania State University, 2003.

Pesendorfer, M., "A Study of Collusion in First-Price Auctions, " (2000) Review of Economic Studies, 67, 381-411.

Porter, R. and J. D. Zona, "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy* 101, 1993.

Porter, R. and J. D. Zona, "Ohio School Milk Markets: An Analysis of Bidding," *RAND Journal of Economics* 30, 1999.

Schuster, E. and M. Niccolucci, "Sealed-Bid Versus Oral-Auction Timber Offerings: Implications of Imperfect Data," *Canadian Journal of Forest Research*, 1994, 24, pp. 87–91.

Stone, J. and D. Rideout, "Does Sealed Bidding Promote Competition in the Auctioning of Public Timber?" *Journal of Forest Economics*, 1997, 3(2), pp. 133–142.

United States House of Representatives, "Timber Sale Methods and Related Practices," Hearings before the Subcommittee on Forests, Committee on Agriculture, on H.R. 5863, February 7-9 and April 5, 1977.

United States Senate, "Timber Sale Bidding Procedures," Hearings before the Subcomittee on Public Lands and Resources, Committee on Energy and Natural Resources, on S. 1360, May 18, 1977.

Vickrey, W., "Counterspeculation, Auctions and Sealed Tenders," *Journal of Finance*, 1961, 16, pp. 8–37.

Table 1A: Summary Statistics for Northern Sales

	Open Auctions				Sealed Auctions			
	Full Sample		Selected		Full Sample		Selected	
N	*787*		*658*		*308*		*300*	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
Auction Outcomes								
Winning Bid ($/mbf)	70.14	52.94	72.78	53.81	80.21	56.25	81.10	56.57
Entrants	4.12	2.46	4.23	2.45	4.53	2.84	4.57	2.86
# Loggers Entering	2.62	2.40	2.84	2.39	3.36	2.58	3.42	2.59
# Mills Entering	1.50	1.65	1.40	1.66	1.17	1.66	1.14	1.66
Fraction Loggers Entering	0.61	0.39	0.65	0.38	0.76	0.32	0.77	0.32
Logger Wins Auction	0.56	0.50	0.60	0.49	0.67	0.47	0.69	0.46
Appraisal Variables								
Volume of timber (hundred mbf)	28.93	39.64	21.95	33.71	15.02	26.97	12.88	22.51
Reserve Price ($/mbf)	26.22	26.72	27.45	27.72	28.46	24.24	28.68	24.38
Selling Value ($/mbf)	196.04	168.41	196.02	169.11	202.59	166.07	201.80	166.66
Road Construction ($/mbf)	6.36	9.84	4.91	9.07	3.11	7.77	2.83	7.54
No Road Construction	0.58	0.49	0.66	0.47	0.78	0.42	0.79	0.41
Logging Costs ($/mbf)	84.66	63.64	82.91	63.77	83.55	62.81	82.51	63.25
Manufacturing Costs ($/mbf)	114.59	84.04	112.93	84.71	117.79	85.57	116.75	86.40
Sale Characteristics								
Contract Length (months)	24.78	17.38	22.19	16.35	18.12	14.79	17.03	13.11
Species Herfindal	0.60	0.27	0.59	0.28	0.58	0.27	0.58	0.27
Density of Timber (hmbf/acres)	0.07	0.06	0.07	0.06	0.08	0.07	0.08	0.07
Salvage Sale	0.37	0.48	0.37	0.48	0.39	0.49	0.40	0.49
Scale Sale	0.44	0.50	0.42	0.49	0.41	0.49	0.40	0.49
Quarter of Sale	2.39	1.00	2.39	1.01	2.42	1.01	2.42	1.01
Year of Sale	86.08	2.31	86.07	2.38	85.75	2.52	85.76	2.55
Housing Starts	1580.62	237.95	1572.33	235.52	1559.18	261.09	1553.84	261.71
Potential Competition								
Logging companies in county	43.86	21.22	42.15	21.67	40.05	22.22	40.36	22.35
Sawmills in County	8.66	4.45	8.42	4.56	7.60	4.47	7.45	4.30
Active Loggers (active in District in prior 12 months)	30.97	24.83	30.19	24.22	25.83	17.62	26.19	17.69
Active Manufacturers (active in District in prior 12 months)	11.02	9.01	11.50	9.26	12.33	10.30	12.54	10.34

Table 1B: Summary Statistics for California Sales

	Open Auctions				Sealed Auctions			
	Full Sample		Selected		Full Sample		Selected	
N	*1188*		*212*		*694*		*269*	
	Mean	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
Auction Outcomes								
Winning Bid ($/mbf)	108.62	165.23	118.95	103.51	93.25	71.80	92.09	74.24
Entrants	4.13	2.32	4.23	2.41	3.85	2.59	4.40	2.68
# Loggers Entering	1.15	1.56	2.12	2.09	2.86	2.25	3.02	2.35
# Mills Entering	2.98	1.81	2.11	1.90	0.99	1.43	1.38	1.58
Fraction Loggers Entering	0.24	0.28	0.50	0.37	0.77	0.31	0.70	0.32
Logger Wins Auction	0.17	0.38	0.43	0.50	0.73	0.45	0.62	0.49
Appraisal Variables								
Volume of timber (hundred mbf)	63.63	45.60	19.85	20.00	7.39	13.38	10.46	10.36
Reserve Price ($/mbf)	41.96	38.02	49.68	46.54	42.56	39.84	37.32	37.09
Selling Value ($/mbf)	278.86	85.30	246.80	131.93	234.49	268.00	247.68	118.60
Road Construction ($/mbf)	10.66	12.95	4.71	11.44	1.08	4.33	2.04	5.89
No Road Construction	0.26	0.44	0.67	0.47	0.90	0.29	0.83	0.38
Logging Costs ($/mbf)	112.85	40.48	96.24	55.24	89.15	56.32	103.47	52.70
Manufacturing Costs ($/mbf)	127.41	34.47	109.20	54.36	100.97	61.85	114.06	52.95
Sale Characteristics								
Contract Length (months)	28.68	14.35	16.37	9.75	10.01	6.62	12.51	6.16
Species Herfindal	0.54	0.23	0.59	0.25	0.60	0.24	0.58	0.24
Density of Timber (hmbf/acres)	0.10	0.12	0.09	0.10	0.16	1.82	0.11	0.15
Salvage Sale	0.14	0.35	0.25	0.43	0.36	0.48	0.26	0.44
Scale Sale	0.95	0.21	0.86	0.35	0.67	0.47	0.82	0.38
Quarter of Sale	2.35	1.00	2.55	0.95	2.71	0.88	2.65	0.93
Year of Sale	85.32	2.14	85.62	2.42	85.59	2.30	85.01	2.15
Housing Starts	1587.06	251.78	1528.56	260.22	1558.48	249.87	1581.44	264.07
Potential Competition								
Logging companies in county	23.22	18.65	22.32	17.56	20.39	17.35	23.06	19.84
Sawmills in County	6.65	6.50	6.14	5.55	6.05	6.01	7.04	7.73
Active Loggers (active in Forest in prior 12 months)	57.65	32.79	60.23	31.55	54.37	30.31	57.96	28.28
Active Manufacturers (active in Forest in prior 12 months)	47.39	27.81	48.96	26.17	44.48	27.08	46.48	26.25

Table 2: Choice of Sale Method

Dependent Variable: Dummy if auction is sealed bid (Logit regression)				
	(1)		(2)	
	Northern		**California**	
	coefficient	**s.e.**	**coefficient**	**s.e.**
Appraisal Controls				
Ln(Reserve Price)	0.006	(0.115)	0.192	(0.180)
Ln(Selling Value)	-0.049	(0.060)	0.196	(0.593)
Ln(Logging Costs)	-0.143	(0.428)	0.302	(0.545)
Ln(Manufacturing Costs)	0.190	(0.426)	-0.646	(0.499)
Ln(Road Costs)	-0.056	(0.208)	-0.025	(0.219)
No Road Construct. (Dummy)	0.455	(0.555)	0.473	(0.565)
Other Sale Characteristics				
ln(Contract Length/volume)	-0.094	(0.254)	0.005	(0.385)
Species Herfindal	-0.735	(0.396)	-0.005	(0.473)
Density of Timber (hmbf/acres)	-1.645	(1.248)	0.162	(0.324)
Salvage Sale (Dummy)	0.167	(0.183)	-0.134	(0.284)
Scale Sale (Dummy)	0.373	(0.195)	-1.509	(0.346)
ln(Monthly US House Starts)	-1.415	(1.049)	-5.965	(1.534)
Volume Controls (Dummy Variables):				
Volume: 1.5-3 hundred mbf	0.072	(0.339)	-1.394	(0.682)
Volume: 3-5	-0.236	(0.378)	-1.611	(0.697)
Volume: 5-8	-0.172	(0.404)	-1.790	(0.747)
Volume: 8-12	-0.754	(0.445)	-2.902	(0.783)
Volume: 12-20	-0.690	(0.478)	-3.632	(0.830)
Volume: 20-40	-1.144	(0.524)	-7.229	(0.924)
Volume: 40-65	-1.785	(0.632)	-8.615	(1.011)
Volume: 65-90	-1.594	(0.723)	-8.320	(1.052)
Volume: 90+	-2.081	(0.705)	-10.013	(1.393)
Potential Competition				
ln(Loggers in County)	-0.276	(0.235)	-0.866	(0.329)
ln(Sawmills in County)	-0.336	(0.296)	0.355	(0.356)
ln(Active Loggers)	-0.058	(0.133)	-0.004	(0.291)
ln(Active Manufacturers)	-0.084	(0.151)	0.234	(0.339)
Constant	11.979	(7.694)	49.668	(11.012)
Additional Controls (Dummy Variables)				
Chi-Squared Statistics (p-value in parenthesis)				
Years	6.25	(0.619)	58.30	(0.000)
Quarters	2.08	(0.556)	0.76	(0.860)
Species	12.14	(0.205)	14.58	(0.006)
Location	78.71	(0.000)	144.09	(0.000)
	N=1095		N=1882	
	LR chi2 (57)	220.11	LR chi2 (50)	1808.59
	P-value	0.000	P-value	0.000
	Pseudo-R2	0.1692	Pseudo-R2	0.7299

Table 3A: Effect of Auction Method on Sale Outcomes (Northern Sales)

(N= 958 Sales)

Panel A: Regression Estimates

Dependent Variable:	ln(Logger Entry)	ln(Mill Entry)	Loggers/Entrants	Logger Wins	ln(Price)	ln(Price)*
No Interactions Between Sealed and Covariates						
Sealed Bid Effect	0.104	-0.017	.056	0.044	0.125	0.076
	(0.037)**	(0.032)	(0.016)***	(0.028)	(0.039)***	(0.032)**
Includes Interactions Between Sealed and All Covariates						
Sealed Bid Effect on Sample	0.105	0.004	.045	0.034	0.139	0.067
	(0.037)**	(0.033)	(0.015)**	(0.028)	(0.041)***	(0.032)*

See Appendix Tables 1A and 2A for full set of controls and coefficients for the no-interaction specifications. Robust standard errors.

Panel B: Matching Estimates

Dependent Variable:	ln(Logger Entry)	ln(Mill Entry)	Loggers/Entrants	Logger Wins	ln(Price)	ln(Price)*
Sealed Bid Effect on Sample	0.158	-0.036	0.079	0.074	0.179	0.133
	(0.043)***	(0.041)	(0.019)***	(0.032)*	(0.052)**	(0 .048)**

Number of matches = 4 using same controls as Panel A and the estimated propensity score. Robust standard errors (using 4 matches).

Note: specification includes number of entering mills and loggers in addition to sale controls.

Table 3B: Effect of Auction Method on Sale Outcomes (California Sales)

(N= 481 Sales)

Panel A: Regression Estimates

Dependent Variable:	**ln(Logger Entry)**	**ln(Mill Entry)**	**Loggers/Entrants**	**Logger Wins**	**ln(Winning Bid)**	**ln(Winning Bid)***
No Interactions Between Sealed and Covariates						
Sealed Bid Effect	0.131	-0.069	0.087	0.086	0.013	-0.048
	(0.058)*	(0.051)	(0.029)**	(0.046)+	(0.065)	(0.055)
Includes Interactions Between Sealed and All Covariates						
Sealed Bid Effect on Sample	0.120	-0.079	0.084	0.077	0.009	-0.027
	(0.058)*	(0.050)	(0.029)**	(0.046)+	(0.064)	(0.048)

See Appendix Tables 1B and 2B for full set of controls and coefficients for the no-interaction specifications. Robust standard errors.

Panel B: Matching Estimates

Dependent Variable:	**ln(Logger Entry)**	**ln(Mill Entry)**	**Loggers/Entrants**	**Logger Wins**	**ln(Winning Bid)**	**ln(Winning Bid)***
Sealed Bid Effect on Sample	0.181	-0.194	0.152	0.135	-0.048	-0.027
	(0.061)**	(0.053)***	(0.031)***	(0.045)**	(0.076)	(0 .075)

Number of matches = 4 using same controls as Panel A and the estimated propensity score. Robust standard errors (using 4 matches).

Table 4: Bid and Entry Distributions for Sealed Bid Auctions

	(1) Bid Distribution (Weibull)		(2) Logger Entry (Poisson)	
	coefficient	s.e.	coefficient	s.e.
	ln(λ)		***ln(μ)***	
Ln(Reserve Price)	0.43	(0.03)	-0.29	(0.05)
Ln(Selling Value)	-0.01	(0.02)	-0.03	(0.03)
Ln(Manufacturing Costs)	0.39	(0.14)	0.85	(0.17)
Ln(Logging Costs)	-0.39	(0.14)	-0.81	(0.17)
Ln(Road Costs)	0.00	(0.03)	-0.16	(0.04)
Species Herfindal	-0.08	(0.11)	-0.24	(0.15)
Density of Timber (hmbf/acres)	-0.88	(0.31)	-0.91	(0.44)
Salvage Sale (Dummy)	-0.05	(0.05)	-0.02	(0.07)
Scale Sale (Dummy)	-0.07	(0.05)	-0.15	(0.08)
Ln(Volume)	-0.07	(0.03)	-0.24	(0.04)
Kootenai NF (Dummy)	0.12	(0.06)	0.18	(0.09)
Mill (Dummy)	0.26	(0.03)		
Mill Entrants	0.12	(0.02)	0.08	(0.03)
Logger Entrants	0.05	(0.01)		
Potential Logger Entrants			0.01	(0.00)
Constant	2.70	(0.19)	2.20	(0.26)
Poisson parameter and Weibull scale parameter include year dummies				
	ln(p)			
Mill(Dummy)	0.00	(0.06)		
Logger Entrants	0.03	(0.01)		
Mill Entrants	0.06	(0.01)		
Constant	0.92	(0.09)		
	ln(θ)			
Constant	-0.52	(0.13)		
	N=1325		N = 300	
	Wald Chi-sq (22)	879.7	LR Chi-sq (21)	199.3
	P-value	0.000	P-value	0.000
			Pseudo-R2	0.14

Note: Bidding distribution model estimated using all sealed bids in 300 sealed bid auctions.

Table 5: Actual Outcomes vs. Outcomes Predicted by Model

	N	(1) Actual	(2) Predicted (bidding only)		(3) Predicted (entry + bidding)	
Sealed Bid Sales						
Avg. Bid	1370	68.5	70.0	(2.0)	66.2	(1.8)
Avg. Logger Bid	1027	57.6	59.2	(1.8)	54.1	(1.6)
Avg. Mill Bid	343	101.0	102.4	(3.7)	102.4	(3.7)
Avg. Sale Price	300	81.1	83.8	(2.2)	85.0	(2.2)
Avg. Revenue	300	116,207	113,663		117,202	
% Sales won by Loggers	300	68.7	66.0	(0.9)	63.5	(1.0)
Avg. Logger Entry	300	3.42	3.42	N/A	3.42	(0.1)
Open Auction Sales						
Avg. Sale Price (Competition)	658	72.8	79.8	(2.3)	79.0	(2.5)
Avg. Sale Price (Collusion)	658	72.8	52.4	(0.0)	51.5	(1.5)
Avg. Revenue (Competition)	658	156,937	163,478		162,820	
Avg. Revenue (Collusion)	658	156,937	62,621		65,930	
% Sales won by Loggers	658	60.0	54.4	(2.3)	51.5	(2.3)
Avg. Logger Entry	658	2.84	2.84	N/A	2.81	(0.4)

Note: Bootstrap standard errors in parentheses

Table 6: Welfare Effects of Sealed vs. Open Auctions

	(1) Sealed Bid	(2) Open Auction (Competitive)	(3) Difference		(2) Open Auction (Part. Collusion)	(3) Difference	
Predict Bidding Only							
Avg. Sale Price	81.5	80.8	0.6	(0.2)	74.5	7.0	(0.4)
Avg. Sale Revenue	148,533	147,665	868	(326)	125,935	22,599	(1612)
Avg. Sale Surplus	259,025	259,167	-141	(31.0)	259,167	-141	(31.0)
% Sales Won by Loggers	58.6	57.8	0.8	(0.1)	57.8	0.8	(0.1)
Predict Entry & Bidding							
Avg. Sale Price	82.9	80.0	2.9		73.4	9.5	
Avg. Sale Revenue	154,620	146,564	8,056		125,277	29,344	
Avg. Sale Surplus	247,298	249,180	-1,882		249,180	-1882	
% Sales Won by Loggers	57.7	54.6	3.2		54.6	3.2	
Logger Entry	3.28	2.91	0.37		2.91	0.37	

Note: Bootstrap standard errors in parentheses. Standard errors are not yet available for the model with predicted entry and bidding.

Figure 1A
Density of Propensity Score by Auction Format for Idaho and Montana Sales—
Full and Selected Samples

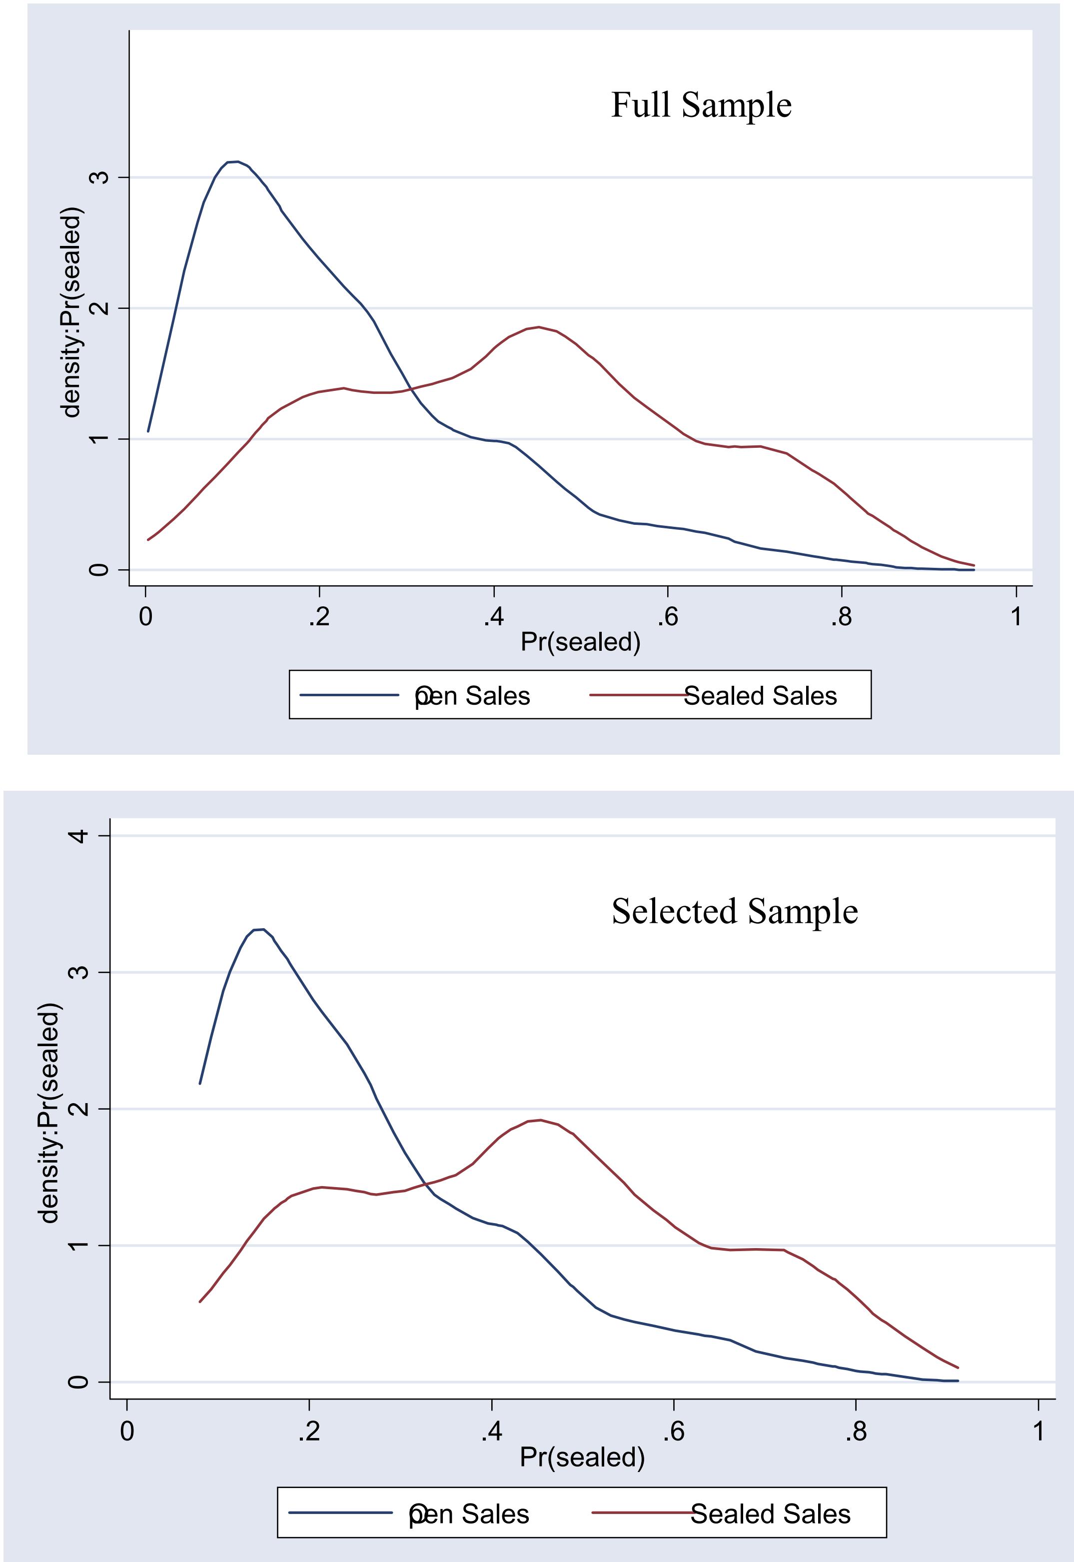

Figure 1B
Density of Propensity Score by Auction Format for California Sales—
Full and Selected Samples

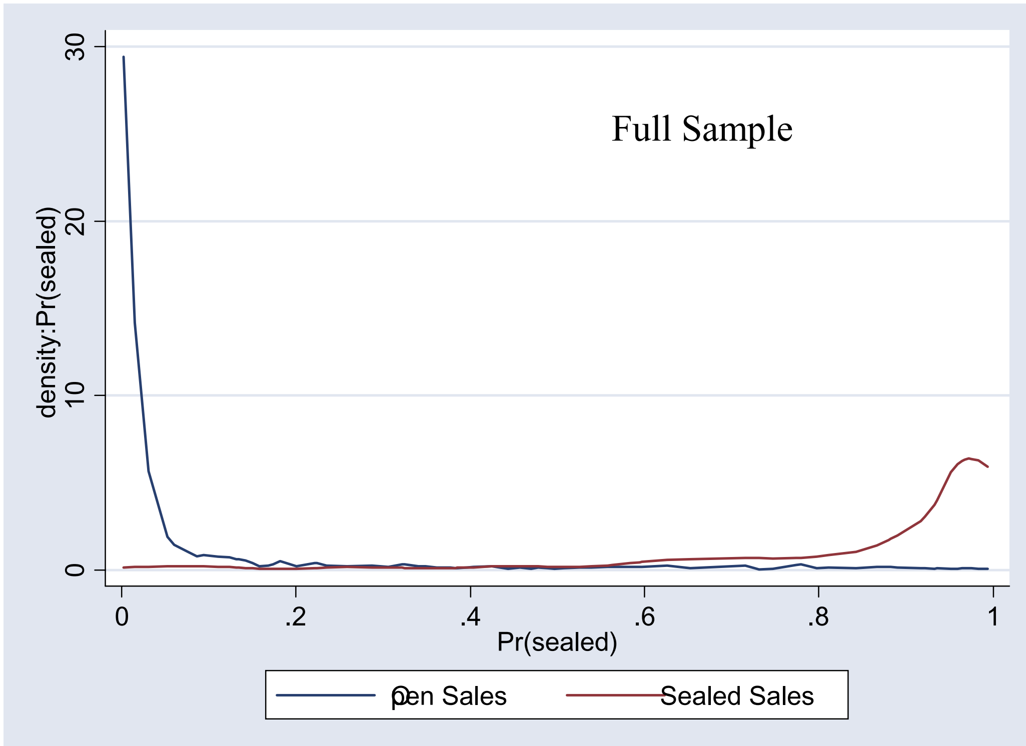

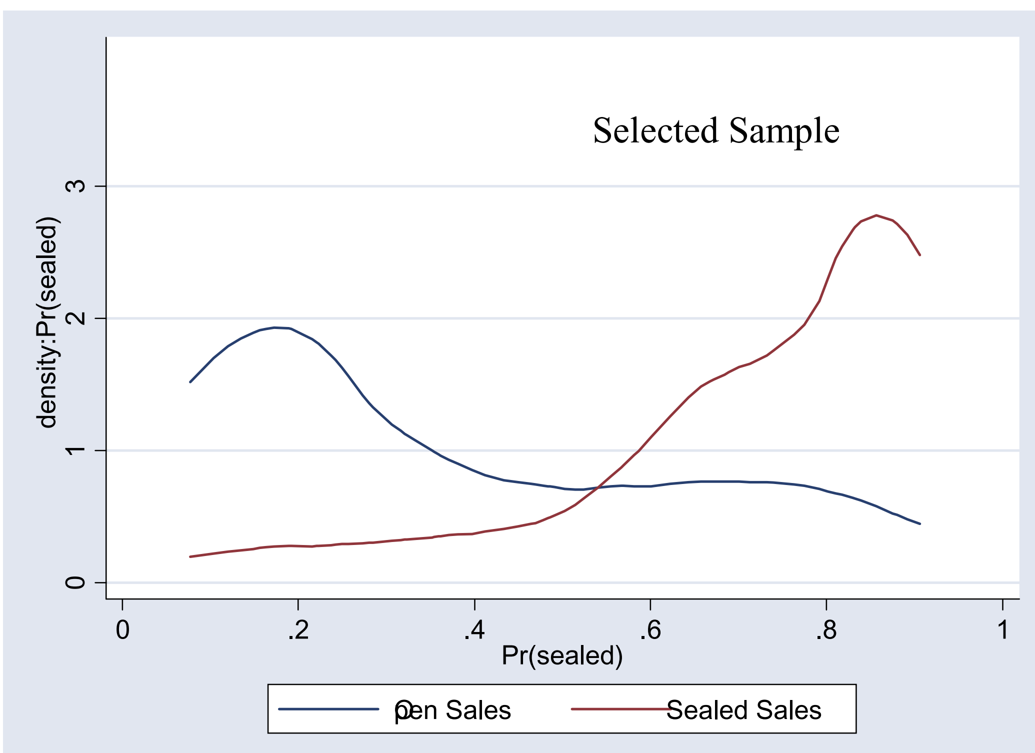

Figure 2: Estimated Bid Functions and Densities of Bidder Valuations for the Case of Two Loggers, Two Mills