

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

Summary

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, shift 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 qualitative 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.

Keywords: Auctions, Timber

JEL Classification: Q23, D44

This paper was presented at the EuroConference on “Auctions and Market Design: Theory, Evidence and Applications”, organised by Fondazione Eni Enrico Mattei and Consip and sponsored by the EU, Rome, September 23-25, 2004.

We thank Phil Haile, Guido Imbens, Richard Levin, Paul Milgrom, and Ilya Segal for helpful suggestions, and Jerry Hausman for advice at the early stages of this project. We are especially grateful to Rob Porter for providing detailed and insightful comments on an earlier draft of the paper. Athey and Levin acknowledge the support of the National Science Foundation.

Address for correspondence:

Jonathan Levin
Department of Economics
Stanford University
Landau Center
Stanford CA94305-6072
USA
Phone: 650 723 5962
Fax: 650 725 5702
E-mail: dlevin@stanford.edu

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, and in asset liquidation sales.

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, have independent and identically distributed values, and bid competitively, 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, entry costs, collusion, common rather than private values, risk-aversion, transaction costs) 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 (1966) argued that open auctions generated less revenue. In 1976,

¹This 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; both areas used a mix of formats during our sample period, 1982-1990.

The theoretical component of our analysis highlights 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. Second, we explicitly model participation by making it costly to acquire information and bid in the auction. Modeling participation adds realism, and more importantly gives rise to new testable hypotheses about entry patterns.²

Our baseline model assumes that firms behave competitively. Under this assumption, auction format has no effect on entry, allocation and revenue when bidders are homogenous. With heterogeneous bidders, however, sealed bidding promotes entry by weaker bidders and can discourage entry by stronger bidders. Sealed bidding also shifts the allocation toward weaker bidders.

To see why sealed bidding favors weaker bidders, 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 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 discourages the entry of strong bidders. We observe, however, that only weak bidder entry is likely to be affected if bidders have similar costs of entry.

²Maskin 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. See also Arozamena and Cantillon (2004) for an analysis of ex ante investment by heterogeneous bidders.

The competitive theory does not generate unambiguous predictions about revenue. Existing examples suggest that with a fixed set of heterogeneous bidders, revenue is often (but not always) higher with sealed bidding. Endogenous entry generates an additional complication because participation varies with the auction format. A revenue comparison, therefore, depends on all the primitives of the model: the bidders' value distributions together with entry costs. Consequently one of our goals 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. The first part examines the qualitative predictions of our theory and quantifies the effect of auction format on observed outcomes. The second part 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 under this assumption, 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 and efficiency suggested by theory. Throughout, 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.

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 also find 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 it can account for the quantitative differences. 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, we consider alternatives to the baseline competitive model. 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 to rationalize our findings. 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 by mills 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. We build on the techniques pioneered by Guerre, Perrigne and Vuong (2000) 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 recover 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 we 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. We also 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 in the Northern Forests. Rather, the data appear consistent with a mild degree of cooperative behavior on the part of participating mills. In contrast, the competitive bidding model appears to fit the open auction prices in California relatively well.

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 increases revenue by roughly 5% relative to a competitive open auction, at minimal cost to social surplus. Strikingly, 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 20%). This suggests that bidder competitiveness merits considerable attention in the choice of auction format.

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.³ 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, Schuster and Niccolucci (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 Schuster and Niccolucci’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 take a somewhat different perspective, we have drawn on their work to select our data sample.

Our work also relates to the empirical literature on collusion at auctions. A variety

³Indeed, most analyses of auctions assume that bidders are symmetric. A few notable exceptions study asymmetries in auctions with fixed participation, including Bajari (1997), Bramman and Froeb (2000), Pesendorfer (2000), Hortacsu (2002), Jofre-Benet and Pesendorfer (2003), and Brendstrup and Paarsch (2003a, 2003b).

of approaches have been suggested to assess whether bidding data are consistent with models of competition or collusion.⁴ 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 as a benchmark from which to evaluate the competitiveness of behavior under an alternative set of rules.

Finally, the last part of this paper shares features with the empirical literature that uses entry decisions to recover estimates of firms' profit functions (Bresnahan and Reiss, 1987; Berry, 1992). This literature uses entry decisions to draw inferences about profit functions relative to a normalized distribution of entry costs, as a function of market-specific covariates. In contrast, we first estimate post-entry profits from firms' pricing decisions (i.e. their bids), and use entry decisions only to recover the sunk costs of participation. This approach allows us to fully recover the parameters of our model in dollar terms.

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 is a set N of 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

⁴Examples 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; they argue that collusion provides a better fit than competition.

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 $[\underline{k}, \bar{k}]$, and each bidder i 's value as a draw from a distribution F_i with support $[\underline{v} = r, \bar{v}_i]$.⁵ Anticipating our empirical analysis, we allow for two kinds of bidders. Bidders $1, \dots, N_L$ are *Loggers* and have value distribution F_L , while bidders $N_L + 1, \dots, N_L + N_M$ are *Mills* and have value distribution F_M . We assume that F_M stochastically dominates F_L according to a hazard rate order, 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 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.⁶

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

⁵The assumption that the reserve price equals the lowest possible value is easily relaxed.

⁶This 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).

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:

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

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:

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

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 $K = (K_1, \dots, K_n)$, bidder i 's expected profit from entry is:

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

where $\Pr[n \mid K, i \in n]$ 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:

$$K_i^s = \Pi_i^s(K^s). \quad (4)$$

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 , despite the fact that (ii) mills shade their bids more than loggers: $b_M(v; n) \leq b_L(v; n)$ for all v .*

The first part of the Proposition states that mills will tend to submit higher bids than loggers. The second part states that mills shade their bids more than 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:

$$\pi_i^o(v_i; n) : \max_{b \geq r} (v_i - \mathbb{E}[\max\{v_{-i}, r\} | v_j \leq b \forall j \in n \setminus i]) \prod_{j \in n \setminus i} G_j(b; n). \quad (5)$$

where $G_j(b; n) = F_j(b)$ is the probability that j bids less than 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:

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

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

$$K_i^o = \Pi_i^o(K^o). \quad (6)$$

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

In equilibrium, mills bid more than loggers. 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 and use identical strategies in equilibrium. Here, an extension of the revenue equivalence theorem implies that the two formats 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 (i) the same expected entry, and (ii) the same expected revenue.*

Revenue equivalence breaks down if bidders are heterogeneous. To analyze this case, we exploit the relationship between a bidder's equilibrium profits and his probability of winning. Given a value v and a set of participants n , bidder i 's expected profit:

$$\pi_i(v; n) = \int_v^v \Pr[i \text{ wins} \mid v_i = x; n] dx. \quad (7)$$

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: (i) loggers are less likely to enter; (ii) mills are more likely to enter; (iii) it is less likely a logger will win.*⁷

⁷The statement of the result is complicated slightly by the fact that there may be several type-symmetric entry equilibria for each auction format. If both formats have a unique entry equilibrium, loggers necessarily enter and win more with a sealed format.

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 strategy profile is an entry equilibrium of the open auction, but every sealed auction equilibrium is inefficient.*

As noted earlier, there is no general theoretical comparison for expected revenue. Existing examples suggest that when participation is fixed, sealed bid auctions often, but not always, result in higher revenue (Maskin and Riley, 2000; Li and Riley, 1999). In principle, endogenous entry could tip the revenue comparison either 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 (Mead, 1966; 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. Loggers simply bid up to their value. We maintain the assumption that bidders make independent participation decisions, so mills anticipate colluding with other participating mills, but do not coordinate entry.⁸

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 only the high-valued mill is relevant in this regard. Nevertheless, collusion gives mills a greater incentive to participate, and this in turn can crowd out logger participation.

⁸There are forms of collusion, such as bid rotation, that involve coordinated entry. 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 for the vast majority of pairs negative correlation can be rejected.

Proposition 6 (*Collusion*) *For any type-symmetric entry equilibrium of the open auction, there is a type-symmetric collusive equilibrium in which: (i) Loggers are less likely to enter; (ii) Mills are more likely to enter; (iii) 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 (i)-(iii) hold.*

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 factors omitted from the model: common values and bidder risk-aversion.

Concentrated versus Dispersed Entry Costs

Our model assumes that bidders differ in their entry costs. In principle the distribution of these costs could be dispersed or highly concentrated; this distinction is relevant for interpreting the results.

If entry costs are dispersed, every potential bidder will be “marginal” in the sense of having a probability of entry strictly between zero and one. In this case, a change in the auction format that changes all bidders’ expected profits will affect the equilibrium entry behavior of both mills and loggers.

In contrast, if entry costs are concentrated so that all bidders have essentially the same entry cost, mills and loggers cannot both be marginal because conditional on entry mills expect higher profits than loggers. In equilibrium, either mills will be roughly indifferent to entering while no loggers enter (clearly not the appropriate assumption for our data), or loggers will be roughly indifferent while mills always enter.⁹ In the latter case, mill participation will be unaffected by auction format,

⁹In the former case, participating 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 possibility.

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.

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).¹⁰ 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 (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.¹¹

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, 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

¹⁰Athey 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.

¹¹See also Perrigne (2003) for evidence of risk aversion from French timber auctions.

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

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.¹² 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.

Before the auction, the bidders have the opportunity to cruise the tract and prepare bids. For sealed bid sales, the Forest Service records the identity of each bidder and their bid. For open auctions, firms must submit a qualifying bid prior to the sale.¹³ Typically these bids are set to equal the reserve price. The Forest Service records the identity of each qualifying firm, as well as the highest bid each qualifier offers during the auction.

¹²The reserve price is computed according to a formula that uses the cruise estimates of timber value and costs, and adds a fixed margin for profit and risk.

¹³This institutional setting is unusual in that there is a record of all bidders in the open auction, even if not all bidders actively bid at the auction. Clearly, the set of bidders must be accurately observed in open auctions for our entry comparisons of open and sealed auctions to be meaningful.

Once the auction is completed, the winner has a set amount of time – typically 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 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 sale announcement. 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 \$80 in the sealed auctions and \$70 in the open auctions. The number of entering logging companies is also somewhat higher in sealed auctions (3.4 versus 2.6), while the number of entering mills is slightly lower (1.2 versus 1.5). 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 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 Schuster and

Niccolucci (1993) report that the choice of sale format was explicitly randomized for a subset of these sales. 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 using the subsample that Schuster and Niccolucci (1993) identify as randomized, though our estimates are somewhat less precise due to the smaller sample size.¹⁴

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 and 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 and the number of logging firms and sawmills in the county of the sale, as counted by the U.S. Census in the past year. Finally, we construct a measure of “active bidders” for each sale by identifying all firms that bid in the forest district in the prior 300 days. We use the number of active loggers and mills to proxy for the number of potential bidders.¹⁵ 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

¹⁴Within our two forests, we include more districts and years than those Shuster and Nicolluci 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.

¹⁵In 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 firms that appear few times it is sometimes hard to distinguish whether two bidders in distinct sales are really the same firm. This is less of a problem with mills as they generally appear many times. Note that for the California sales, we use the forest rather than the district as a unit of analysis, as there are fewer sales per forest.

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.

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, 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.¹⁶

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, may 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, 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

¹⁶The dropped sales are generally large volume sales in districts that ran few sealed auctions. This criteria leads us to drop a larger fraction of sales in California, where sale method correlates more closely with sale size.

the same classification.¹⁷ Mills also attend more auctions than most loggers, although a few loggers attend frequently.

Our theoretical model assumes that mills tend to have higher willingness to pay than loggers. An implication is that mills should submit higher bids and win disproportionately. To check this, 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%).¹⁸

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

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

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, $N = (N_L, N_M)$ is the number of potential bidders, and ε is unobservable.

A standard point is that to identify the average effect of auction format, denoted $\tau_Y = \mathbb{E}_{X,\varepsilon}[f(1, X, N, \varepsilon) - f(0, X, N, \varepsilon)]$, we require that the unobserved component

¹⁷The employment data are 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 appearances by mills in our data, 1311 are by mills that are large. In contrast, only 467 of 3097 logger appearances are by large firms.

¹⁸Mills also have a higher entry rate than loggers (which is a feature of the most natural equilibria of our model), although our measurement of this suffers from the difficulty noted above of precisely identifying potential bidders. Using our measure of potential bidders, 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.

of the outcome is independent of the auction format conditional on covariates. This 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).¹⁹ It holds for the other sales if the choice of format is based on information from the Forest Service appraisal, or follows some rule based on covariates in our data.²⁰

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

$$Y = \alpha \cdot SEALED + X\beta + N\gamma + \varepsilon. \quad (9)$$

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 and compute its average effect for the sample. A second issue is that we must specify the functional form for the covariates to be included in X . While our results are not very sensitive to the alternatives we have tried, in principle misspecification could lead to bias.²¹

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

¹⁹Otherwise, we could mis-estimate 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 Schuster and Niccolucci (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.

²⁰If 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 characteristics 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.

²¹There 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 volume might lead us to falsely impute a revenue effect of auction method.

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.²² 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:

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

where T^s and T^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 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

²²We 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.

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 sealed bid. These estimates are at best marginally statistically significant. Thus, although our point estimates are not insubstantial, we cannot rule out a fairly small effect of auction format on allocation.

Finally, we turn to revenue. 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 that 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 an 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.²³

C. Evidence from California Forests

While the Northern forests seem particularly well-suited to making a statistical

²³An 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 coefficients were highly sensitive to the particular choice of potential competition measures, none of which are ideal.

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.

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% more loggers at sealed sales, which

translates into an additional 3 loggers participating for every 10 sales. 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 indicates that sealed bid sales 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, but larger, than those in the Northern forests. In the California forests, we estimate that there is roughly a 8-14% greater chance a logger will win with sealed bidding.

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. Moreover, the key assumption generating

these departures from revenue equivalence, that bidders are heterogeneous, also seems consistent with the data.²⁴

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 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 more bidder interest, especially on the part of logging companies. This would help to explain the entry results, though it is not clear to us why forest managers would systematically behave

²⁴Above, 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 to participate and to win conditional on participating.

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.²⁵

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.

5. Structural Estimation and Testing

Our final goal is to assess quantitatively the relationship between our findings and the theory we proposed to account for them. We investigate three related issues. First, we ask whether 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 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.

The key elements of our approach are as follows. We use entry and bidding data from the sealed bid auctions to estimate the parameters of our theoretical model — the value distributions of loggers and mills, and the costs of entry — as functions of observed tract characteristics. To do this, we assume competitive behavior in the sealed bid auction as outlined above. We also allow for unobserved heterogeneity in

²⁵The Forest Service Handbook also instructs forest managers to use sealed bidding if they expect a sale *not* to be competitive.

the underlying values of the tracts. We then use the calibrated model to predict the equilibrium outcome of each sale in our sample and compare the predictions to the actual outcomes. For tracts sold by sealed bidding, this provides a measure of how well our model fits the data. For tracts sold by open auction, the predictions are out-of-sample because the open auctions were not used to estimate the parameters of the model.²⁶ Comparing the predictions to outcomes allows us to assess whether the model accurately accounts for the observed differences across auction formats. It also provides a way to evaluate the competitiveness of open auctions. Finally, we develop a welfare comparison of open and sealed bidding. Paralleling the previous section, we focus on sales in the Northern forests and discuss California sales later.

A. Structural Estimation

Our first step is to use the sealed bid data to estimate the parameters of the theoretical model as a function of tract characteristics. To estimate the value distributions of mills and loggers, we build on the 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. Given the value distributions, we can estimate entry costs using observed entry behavior.

We modify Guerre, Perrigne and Vuong (2000)'s approach in order to account for unobserved heterogeneity in auction characteristics.²⁷ Formally, we let u denote an auction characteristic known to participating bidders but not observed in our data. Let X denote the set of sale characteristics known both to the econometrician

²⁶In fact, Athey and Haile (2002) show that when values are correlated, as in our model of unobserved heterogeneity, underlying value distributions cannot be identified from bidding data in open auctions. Haile and Tamer (2003) point out additional concerns with drawing inferences from losing bids in open auctions.

²⁷We are motivated to include unobserved heterogeneity because there is significant positive correlation among bids in the sealed auctions conditional on observed auction covariates. We also believe it is reasonable to assume that bidders can commonly observe certain features of a tract that make it more or less valuable. An alternative way to rationalize correlation in bids is with an affiliated values model, e.g. where bidders receive a one-dimensional signal and cannot disentangle the common component from their idiosyncratic preferences. Krasnokutskaya (2002) exhibits simulations whereby, relative to assuming independent values and unobserved heterogeneity, using the assumption of affiliated values when estimating values from bidding data leads to larger inferred values, which would increase predicted revenue in open auctions and thereby magnify the revenue difference in our sample.

and the bidders, and let $N = (N_L, N_M)$ represent the number of potential mill and logger entrants. We write the bidders' value distributions, conditional on (X, u, N) , as $F_L(\cdot|X, u, N)$ and $F_M(\cdot|X, u, N)$.

In line with our model, we assume that bidders' values, and hence their bids, are independent conditional on (X, u, N) . As we assume the set $n = (n_L, n_M)$ of participating bidders is observed prior to bidding, we write the equilibrium bid distributions as $G_L(\cdot|X, u, N, n)$ and $G_M(\cdot|X, u, N, n)$. If there is a single bidder, we assume he optimally bids the reserve price, but otherwise we treat the reserve price as non-binding.²⁸

Also in line with our model, we assume that bidders have independent entry costs concentrated around some average cost $K(X, N)$, and that in equilibrium loggers are the marginal participants regardless of auction format. This of course means that the number of potential and actual mill entrants is identical, $n_M = N_M$.²⁹ Note that under these assumptions, we can infer exactly the number of potential mill entrants, rendering our earlier proxy — the number of “active” mills — irrelevant. Finally, we maintain the standard assumption that auctions in our sample are independent of one another.

Estimating the Bid Distributions

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

²⁸See Haile (2001) for a discussion of why Forest Service reserve prices are typically non-binding. A slight drawback to this assumption is that our fitted bid distributions will assign positive (though typically small) probability to bids below the reserve price. We did experiment with modeling bidder values (and hence bids) as being distributed above the reserve price, but found that this model fit the data poorly, possibly because the mechanical formula used to determine the reserve price may not track changes in bidder values over time or across auctions well.

²⁹In principle, it would also be possible to estimate a model of dispersed entry costs, where both logger and mill entry would be stochastic. Identifying an entire *distribution* of entry costs, however, would additional extrapolation across auctions as well as an instrument that affects values but not entry costs. We judged this too much to ask of our data. Note that our current approach implies that mill entry will not vary with auction format. This is not a problem for the Northern region, but is not a perfect fit for California.

After extensive experimentation, we settled on a specification of Weibull bid distributions with Gamma distributed auction heterogeneity. Thus we assume that for $k = L, M$:

$$G_k(b|X, u, N, n) = 1 - \exp\left(-u \cdot \left(\frac{b}{\lambda_k(X, N, n)}\right)^{p_k(n)}\right). \quad (10)$$

Here $\lambda_k(\cdot)$ is the scale, and $p_k(\cdot)$ the shape, of the Weibull distribution, parametrized as $\ln \lambda_k(X, N, n) = X\beta_X + N\beta_N + n\beta_{n,k} + \beta_{0,k}$ and $\ln p_k(n) = n\gamma_{n,k} + \gamma_{0,k}$.^{30,31} We assume u has a Gamma distribution with unit mean and variance θ , and is independent of X , N , and n .³² We estimate the parameters (β, γ, θ) by maximum likelihood; the likelihood function is written out in the Appendix. The estimates for the Northern region are reported in Table 4A.

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. 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 (a property that can potentially be violated if bidder values are affiliated or have a common value component). Finally, we estimate that u has significant variance, indicating that our modeling of unobserved heterogeneity across auctions is warranted.

One natural question concerns the appropriateness of the Gamma-Weibull func-

³⁰The specification we adopt is more parsimonious than in our earlier regressions. Our results do not seem sensitive to including additional covariates; nevertheless, we opted for parsimony because of the need to make out-of-sample predictions where over-fitting could in principle be a problem.

³¹Specifying how the number of participants should affect the bid distribution is a challenge in two-stage structural estimation of auction models, because there is no easy way to incorporate the theoretical restriction that the value distributions be independent of the number of bidders. Theory does predict that mill behavior could be quite different if there is only a single mill, which motivates us to include a single mill effect in the mill bid distribution. Theory also predicts that the effect of an additional bidder on a given bidder's behavior should be limited as the number of bidders grows. For this reason, use $\min\{n_L, \bar{n}\}$ and $\min\{n_M, \bar{n}\}$ in place of n_L, n_M in our estimates, where $\bar{n} = 5$.

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

tional form. Initial experimentation, and our desire to include a rich set of covariates and unobserved auction heterogeneity, led us toward a parametric approach.³³ In addition, existing non-parametric methods typically propose to estimate bid distributions separately for each set of participants (n_L, n_M) ; we do not have sufficient data to do this. We considered several parametric alternatives, each sharing the property with the Gamma-Weibull that bid i in auction t could be written as $b_{it} = \exp(X_t\beta_X + N_t\beta_N) \cdot \varepsilon_{it}(n)$ with some parametric family for the joint distribution of the residuals $\varepsilon_{it}(n)$. We examined how each alternative matched the observed distribution of logger and mill bids, the within-auction bid correlation, and the observed sealed bid prices.

The Gamma-Weibull form appeared to provide a good fit on all dimensions. To provide a sense of this, Figure 2 plots the distribution of sealed bid residuals in our sample (i.e. the distribution of the $\hat{\varepsilon}_{it}$ s, where $\hat{\varepsilon}_{it} = b_{it} / \exp(X_t\hat{\beta}_X + N_t\hat{\beta}_N)$) next to the distribution predicted by our fitted model. The overall mean of the bid residuals is 26.8; the variance is 15.6; the between-auction variance is 12.4 and the within-auction variance is 9.6. By way of comparison, the fitted model predicts a mean of 26.4, and respective variances of 16.1, 12.7 and 9.3. We provide further evidence on how the model fits prices and logger and mill bids in Table 5, discussed below.

Estimating the Value Distributions

We now turn to recovering the bidders' value distributions. Under the assumption that the observed bids are consistent with equilibrium behavior, each bid must be optimal against the opponents' bid distributions. That is, a bidder's value v_i is related to his observed bid b_i through his first-order condition for optimal bidding:

$$v_i = \phi_i(b_i; X, u, N, n) = b_i + \frac{1}{\sum_{j \in n \setminus i} \frac{g_j(b_i | X, u, N, n)}{G_j(b_i | X, u, N, n)}}. \quad (11)$$

It is straightforward to construct an estimate of ϕ_i given our estimates of G_L and G_M . If all sale characteristics (X, u, N, n) were observed, we would then be able to

³³For models without unobserved heterogeneity, Guerre, Perrigne and Vuong (2000) propose a non-parametric approach to estimate the bid distributions, while Jofre-Bonet and Pesendorfer (2003) use a Beta-Weibull specification. Our approach builds on Krasnokutskaya (2002), who introduces unobserved heterogeneity in a semi-parametric model under that assumption that bidder values are additively or multiplicatively separable in u . Bidder values do not have that property in our model.

infer the bidder value corresponding to each observed bid, and thus recover the value distributions (as in Guerre, Perrigne and Vuong, 2000). As u is unobserved, however, we need to modify the approach. As observed by Krasnokutskaya (2002), we can still recover the distributions $F_L(\cdot|X, u, N)$ and $F_M(\cdot|X, u, N)$ for any value of u from the relationship:³⁴

$$F_k(v|X, u, N) = G_k(\phi_k^{-1}(v; X, u, N, n)|X, u, N, n).$$

Figure 3 plots the density functions for logger and mill values for an auction with average covariates, and $u = 1$, as well as the equilibrium bid functions assuming two mills and two loggers participate in the auction.³⁵ 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.³⁶

It is also possible, by averaging across values of u , to estimate the typical markups built into the sealed bids in our data. We estimate that across mill bids, the median profit margin is 10.4%; for loggers the median profit margin is 9.0%.

Estimating Entry Costs

The remaining parameter of the model is the entry cost, which we recover using the equilibrium entry condition. Recall that when entry costs are concentrated and loggers are the marginal participants as we have assumed, then in equilibrium each logger will be nearly indifferent to participating. In particular, a sealed bid entry equilibrium requires that $K(X, N) \approx \Pi_L^s(X, N)$, where $K(X, N)$ is the average entry cost as a function of observed sale characteristics (assumed to be independent of auction

³⁴A small subtlety here is that our theoretical model implies that the equilibrium bid distribution will have a finite upper bound. The Weibull distribution does not. For this reason, we truncate the very upper tail of the estimated distributions $G_L(\cdot)$ and $G_M(\cdot)$ and work with the truncated distributions. The motivation for this and details of the implementation are described in the Appendix.

³⁵To compute these, we started with the fitted bid distributions $G_L(\cdot|X, u, N, n)$ and $G_M(\cdot|X, u, N, n)$, with $X = \bar{X}$, $N = \bar{N}$, $u = 1$ and $n = (2, 2)$, then used the first-order condition to recover the bid functions $b_k(v|X, u, N, n) = \phi_k^{-1}(v|X, u, N, n)$.

³⁶Note that this finding requires more than an ordering of means or a first-order stochastic dominance ordering; rather, it reflects an ordering of the reverse hazard rates, $g_M/G_M \geq g_L/G_L$.

format) and $\Pi_L^s(X, N)$ is the equilibrium profit a logger expects from entering.³⁷

We write this expected profit as:

$$\Pi_L^s(X, N) = \sum_{n \subset N} \bar{\pi}_L^s(X, N, n) \Pr[n|X, i \in n, N, s]. \quad (12)$$

The first term, $\bar{\pi}_L^s(X, N, n)$, is a logger’s expected profit conditional on sale characteristics and the set of participants. We compute this number from our estimate of the value and bid distributions, integrating out the unobserved auction heterogeneity.

The second term, $\Pr[n|X, i \in n, N, s]$, is the equilibrium probability that a logger entering a sealed bid auction (denoted by s) assigns to the set of participants being n , conditional on sale characteristics (X, N) . To estimate this term, we model the number of entering loggers as a Poisson random variable with mean $\mu(X, N)$, parameterized as $\mu(X, N) = X\alpha_X + N\alpha_N$.³⁸ As before, we use our measure of “active” loggers to proxy for the potential logger entrants. We estimate $\mu(X, N)$ using the entry data from the sealed bid auctions in our sample; the estimates are reported in Table 4. We also know that all potential mill entrants will enter in equilibrium, so $n_M = N_M$. We therefore assume firms face no uncertainty about mill entry.

Putting this all together, we use (12) to obtain the predicted logger profits from a sealed bid auction, $\Pi_L^s(X, N)$, as a function of the characteristics (X, N) . Then, treating each tract in our sample as an (X, N) pair, we impute for each tract an entry cost equal to $K(X, N) = \Pi_L^s(X, N)$. We estimate a median entry cost of \$4695 (s.e. \$1132). 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.³⁹

B. Comparing Predicted and Actual Outcomes

Having estimated the parameters of the theoretical model as functions of observable sale characteristics, we now ask how closely the model’s equilibrium predictions

³⁷Note the slight change of notation from the theoretical model; we now include covariates X, N as an argument of the profit function, and suppress the entry thresholds.

³⁸In 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.

³⁹As a point of comparison, we estimate that across tracts in our sample the median expected mill profit from a sealed bid auction is roughly \$45,000 gross of entry costs.

match the observed outcomes in our data. 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 the open auction outcomes, and in particular, whether assuming some degree of cooperative 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 earlier.

To generate sealed bidding predictions, our estimated Poisson model of logger entry gives the equilibrium distribution of loggers who will participate in a sealed bid auction as a function of tract characteristics. The number of mill entrants is known and not stochastic. We use our estimates of G_L, G_M and the distribution of unobserved heterogeneity to predict bidding behavior conditional on participation. Finally we combine the entry and bidding predictions to predict outcomes conditional only on tract characteristics.

To generate open auction predictions, we observe that conditional on participation, each entrant will bid his value, and the auction price will equal the second highest value. Alternatively, if mills collude, all but the highest value mill drop out immediately, and the remaining bidders behave competitively. These observations allow us to calculate expected prices and profits for a given tract, and any given set of participants, under both the assumption of competitive and collusive behavior, by repeated simulation. Each simulation involves drawing a value of u , then drawing a value for each participant from either $F_L(\cdot|X, u, N, n)$ or $F_M(\cdot|X, u, N, n)$, and finally calculating the auction price, profits and surplus.

This gives predicted open auction outcomes for each tract conditional on any hypothetical set of participants. To predict open auction entry, we continue to treat mill entry as known and not stochastic. We assume, as we did earlier, that the equilibrium distribution of logger entrants can be approximated as a Poisson distribution. For each tract, we find the Poisson parameter for which the expected logger profits from entering just equal the entry cost. This yields a prediction of the equilibrium distribution of logger entrants, which we combine with our bidding predictions, to generate predicted outcomes as a function of observed tract characteristics.

Table 5 reports the actual average outcomes in our sample and the average out-

comes predicted by our parameterized model.⁴⁰ For the tracts sold by sealed bid auction, we closely predict the average bids of loggers and mills. In reality, the average bids are \$57.6 and \$101.0. The model predicts averages of \$56.6 and \$101.5 per mbf unconditionally (i.e. given just sale characteristics), and \$55.8 and \$101.5 conditional on the set of participating bidders. We also closely predict the average auction prices and average sale revenue. The model somewhat under-predicts the fraction of sales that loggers win — both the unconditional prediction of 64.5% and the prediction of 67.0% 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 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.

We start by comparing the model’s predictions for entry and allocation to the actual outcomes for the tracts sold by open auction. Strikingly, the model predicts a level of logger entry that is very close to the actual level (2.89 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 53.0% of the sales, or 55.2% conditional on realized participation, while in reality they win 60.0%). Note, however, that despite under-predicting logger purchases for both sale formats, the model accurately captures the difference across the open and sealed bid sales (the model’s prediction is 11.5% versus 8.7% in reality).

A key point 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 distinguish between a range of behavioral assumptions — from

⁴⁰We 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.

competitive behavior by the mills, to perfectly collusive behavior, to any intermediate degree of collusion that involves the highest-valued mill bidding up to his value — it is necessary to focus on prices.

The numbers in Table 5 indicate that the observed prices in the open auctions lie between the competitive and fully collusive prices predicted by the model. The competitive model predicts an average price of \$79.7, or \$80.7 conditional on realized entry. It predicts an average price of \$51.7 per mbf if mills fully collude. In reality, the average sale price across open auctions is \$72.8 per mbf. Even accounting for sampling error, we reject both the competitive and collusive models at conventional confidence levels. Thus the assumption of mildly cooperative behavior on the part of participating mills seems to provide a better match than either the competitive or fully collusive extremes.⁴¹ It is worth noting that this conclusion is not sensitive to our assumption that the sealed bid auctions are competitive. If we assumed a degree of collusion in the sealed bid auctions, we would infer a higher distribution of bidder values from the data. This would reinforce the finding that open auctions appear less than perfectly competitive.

To summarize, it appears that the theoretical model developed in Section 2 and estimated using the sealed bid data does a reasonable job of explaining the differences in outcomes across auction formats we observe in the data. The best fit comes under the assumption that mill behavior in open auctions is mildly cooperative.

C. Evidence from California Forests

To further assess the model’s ability to explain our empirical findings, we repeated our analysis on the California sales. Here the conclusions regarding bidder competitiveness are rather different, with an assumption of competitive behavior providing a plausible fit to the data.

Table 4B reports our estimates of the entry and bid distributions for the California sealed bid auctions. As in the Northern region, we find substantial differences between loggers and mills, and significant unobserved heterogeneity across auctions. The Gamma-Weibull bid distribution again appears to provide a good fit to the data. Figure 2B plots the density of bid residuals for the California sealed bid sales next to

⁴¹A possibility is that there is collusion at a small fraction of the sales. We should note, however, that when we looked at the open auctions for which the predicted price is substantially above the actual price, we did not find any obvious pattern.

the distribution predicted by the fitted model, showing a reasonable match.⁴² Table 5B compares the actual outcomes of these auctions to those predicted by the fitted model. The model’s predictions of average logger and mill bids, average sale prices and the fraction of sales won by loggers match the data fairly closely, while the model somewhat over-predicts sale revenue relative to the realized outcome.

When we move to making out-of-sample predictions about the open auctions in California, we find that in contrast to the Northern region, the observed prices can be described reasonably well under the assumption that firms bid competitively. The average sale price in the California open auctions was \$119.0. As reported in Table 5B, our fitted model predicts an average price of \$120.7 conditional on realized entry, and \$115.1 when we predict entry as well as bidding. The model also predicts logger entry and the fraction of sales won by loggers with some accuracy. As with the sealed bid sales, we somewhat over-predict sale revenue; the revenue difference across formats is relatively close to the actual difference.

D. Quantifying the Trade-offs in Auction Design

So far we have tried to assess if our theoretical model 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 we 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 in both entry and allocation, 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 conduct a welfare comparison, we use our estimates of the primitives to compute the predicted outcome of both an open auction and a sealed bid auction for each tract in our sample. For each tract, and each auction format, we compute the expected entry, the expected price and revenue, the probability that a logger will win, and the expected surplus (the value of the winning bidder net of entry costs sunk by

⁴²In terms of the covariance structure of the bid residuals, the observed mean in our sample is 11.2; the variance is 6.9; the between-auction variance is 4.8; and the within-auction variance is 4.2. Our fitted distribution predicts a mean of 11.0, and corresponding variances of 6.2, 4.9 and 3.7.

all the bidders). For the open auction format, we consider two alternative specifications of mill behavior: 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 in the Northern region).

Our comparisons are reported in Tables 6A and 6B. 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 small, despite substantial asymmetries among bidder types. Sealed bidding would generate more revenue, but the revenue gain is only \$651 per sale in the Northern region and \$1018 in California. Sealed bidding also increases the probability that sales are won by loggers, but the average increase in probability is less than 1%. Finally, the efficiency benefit to using an open auction format is also quite small, only \$100 per sale in the Northern region and \$45 per sale in California.

These differences increase when we account for the fact that bidder participation will vary systematically with auction format, though the estimates are also less precise. 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 3-6 more loggers for every 10 sales. One effect of this additional entry is to generate a more substantial difference in the fraction of sales won by loggers —we predict that loggers would win 3-4% more sales with sealed bidding. A second effect is to increase the revenue advantage of sealed bidding to roughly \$5300 (3%) for the average sale in the Northern region and \$26,000 (13%) in California. Our estimate of the social surplus differential is quite noisy, so much so that our point estimates indicate higher social surplus from sealed bidding, despite the fact that Proposition 5 shows that sealed bidding is less efficient.⁴³

⁴³The reason it is even possible to generate a positive point estimate here is that in practice we estimate separate value distributions for each possible configuration of entrants (n_L, n_M) , and these estimates are not precisely the same. As noted earlier, this is an issue anytime one uses current two-stage auction estimation methods; it becomes visible here because in modeling stochastic logger

As a practical matter, however, the model suggests that these differences are dwarfed by the potential effects of bidder collusion. In the Northern region, 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 \$28,000 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 in this region, the revenue benefits of sealed bidding clearly seem to be the most quantitatively significant welfare consequence of the choice of auction method. In contrast, in California where competitive behavior seems consistent with the observed outcomes, a welfare trade-off between sale formats appears to hinge on relatively small differences.

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 empirical finding is that sealed bid auctions attract more small bidders, shift the allocation toward these bidders, and in some 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 approach to structural estimation in this setting is novel in several ways. First, motivated by a desire to match key features of the application, we use an approach that incorporates several elements (heterogeneous bidders, unobserved auction heterogeneity, and a model of bidder participation) that have previously received attention in isolation. Second, we exploit the variation in auction format to assess the competitiveness of the open auction format. By relying only on data from sealed bid auctions to estimate our primitives, we are able to make out-of-sample predictions for open auctions that can be compared to actual outcomes.

entry we need to take expectations that average over possible numbers of logger entrants, where the weights on different realizations of n_L vary across auction formats. Note that we could take the approach of averaging our value distribution estimates to create a pooled estimate, but this would have the drawback that for any given set of participants, our pooled value distribution estimate would not correspond through the first order condition to the estimated bid distribution.

Even though the role of asymmetries in determining optimal auction design have received a fair amount of attention in the theoretical literature, our results show that with fixed participation, the choice of auction format has little impact even with substantial asymmetries among bidders. When participation is endogenous, we see that sealed bidding favors the small or weak bidders in both entry and allocation, and differences across auction formats are magnified. Finally, our 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 I: 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 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 $\Pi_i(K) = K_i$ for all i . So establishing a type-symmetric entry equilibrium amounts to finding a type-symmetric fixed point of $\Pi = (\Pi_1, \dots, \Pi_n)$. Let $\mathcal{K} = \{(K \in [0, \bar{K}]^n : K_1 = \dots = K_L, K_{L+1} = \dots = K_N)\}$ denote the space of type-symmetric entry thresholds, where $\bar{K} \geq \bar{k}$ is large enough so that no bidder's profits could exceed it. Now, $\Pi : \mathcal{K} \rightarrow \mathcal{K}$ and 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, 2004) 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 $K = (K, \dots, K)$. It follows that both auctions have a unique symmetric entry equilibrium that solves $\Pi_i(K, \dots, K) = K$. The results follow directly.

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

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


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.



Graphical Depiction of Entry Equilibria

To proceed, we characterize type-symmetric entry equilibria of the open auction. Consider the space $\{(K_L, K_M) : K_i \in [0, \bar{K}]\}$ of type-symmetric entry thresholds. Let

\mathcal{L}_L denote the locus of points (K_L, K_M) for which $\Pi_L^o(K) - K_L = 0$, 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 downward sloping because $\Pi_i^o(\cdot)$ is decreasing in K_j . So above \mathcal{L}_i , $\Pi_i(K) - K_i < 0$ while below \mathcal{L}_i , $\Pi_i(K) - K_i > 0$. Moreover, if $K_M = \bar{K}$ then $\Pi_M(K) - K_M < 0$, so \mathcal{L}_M intersects the x -axis at some $\bar{K}_M < \bar{K}$. Moreover, \mathcal{L}_L must be above the x -axis \bar{K}_M because $\Pi_L^o(K) > 0$ if $K_L = 0$. 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 $K_L^o \leq K_L$.

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:

$$\Pi_M(K^s) - K_M^s \geq 0 \quad \text{and} \quad \Pi_L(K^s) - K_L^s \leq 0.$$

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. A socially efficient entry and bidding profile maximizes social surplus. Given participation, efficient bidding means each bidder uses an identical increasing bid strategy, so the bidder with the highest value wins. The sealed bid auction already fails to be efficient on these ground, but the open auction has this property. We now show that open auction also involves efficient entry.

To this end, observe that given a set of participants n , the efficient (and open auction) surplus $s(n)$ equals the expected highest value. Therefore, for any bidder $i \in n$,

$$s(n) - s(n \setminus \{i\}) = \Pr[v_i \geq v_j \forall j \in n \setminus \{i\}] \cdot \mathbb{E} \left[v_i - \max_{j \in n \setminus \{i\}} v_j \mid v_i \geq v_j \forall j \in n \setminus \{i\} \right].$$

This just equals i 's expected profit, denoted $\pi_i(n)$. If $i \notin n$, then clearly $\pi_i(n) = 0$.

Now, let $S(K)$ denote the surplus given entry thresholds $K = (K_1, \dots, K_N)$:

$$S(K) = \sum_n s(n) \Pr[n|K] - \sum_i EC_i(K_i),$$

where $EC_i(K_i) = \mathbb{E}[k_i | k_i \leq K_i] H(K_i)$ is the expected entry cost sunk by i . Bidder

i 's expected profit (net of entry costs) is:

$$\begin{aligned}
u_i(K) &= \Pi_i(K) - EC_i(K_i) \\
&= \sum_n \pi_i(n) \Pr[n|K] - EC_i(K_i) \\
&= \sum_n (s(n) - s(n \setminus \{i\})) \Pr[n|K] - EC_i(K_i) \\
&= S(K) - S(K_i = \underline{k}, K_{-i}).
\end{aligned}$$

Therefore $du_i/dK_i = dS/dK_i$. It follows that given K_{-i} the choice of K_i that maximizes social surplus also maximizes i 's net profits, so the socially efficient entry profile is an entry equilibrium. The stated Proposition, however, involves a restriction to type-symmetric profiles. To see why that poses no problem, let $S^{TS}(K_L, K_M)$ denote the social surplus from a type-symmetric profile (K_L, K_M) (so $S^{TS}(K_L, K_M)$ equals $S(K_L, \dots, K_L, K_M, \dots, K_M)$). And assume K^* is an efficient type-symmetric profile:

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

If K^* is interior (the proof is similar for boundaries), $dS^{TS}/dK_L(K^*) = dS^{TS}/dK_M(K^*) = 0$. Moreover, $dS^{TS}/dK_L = \sum_{i=1}^L dS/dK_i$ and likewise for mills. Because S is symmetric in logger thresholds and also in mill thresholds, it follows that $dS/dK_i(K^*) = 0$ for all loggers and all mills. Thus $du_i/dK_i(K^*) = 0$ for all i , so K^* is 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:

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

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 is sufficiently large, 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.*

Appendix II: Omitted Details of the Structural Model.

A. The Likelihood Function

A useful property of Gamma-Weibull models is that the unobserved heterogeneity can be integrated out analytically. This leads to the following log-likelihood for auction t :

$$\begin{aligned} \ln L_t &= (n_{Lt} + n_{Mt}) \ln \theta + \ln \Gamma \left(\frac{1}{\theta} + n_{Lt} + n_{Mt} \right) - \ln \Gamma \left(\frac{1}{\theta} \right) \\ &+ \sum_{i=1}^{n_{Lt}+n_{Mt}} \ln \left(p_{it} \lambda_{it} \left(\frac{b_{it}}{\lambda_{it}} \right)^{p_{it}-1} \right) + \left(\frac{1}{\theta} + n_{Lt} + n_{Mt} \right) \ln \left(1 + \theta \sum_{i=1}^{n_{Lt}+n_{Mt}} \left(\frac{b_{it}}{\lambda_{it}} \right)^{p_{it}} \right). \end{aligned}$$

Here θ is the Gamma variance, $b_{1t}, \dots, b_{(n_{Lt}+n_{Mt})t}$ are the observed bids in auction t , and λ_{it}, p_{it} are the Weibull parameters for bidder i in auction t . As defined in the text, these are functions of (X_t, N_t, n_t) , the unknown parameter vectors β and γ , and bidder i 's type — logger or mill.

B. Truncating the Bid Distributions

Our independent private values model predicts that the equilibrium bid distributions will have finite support. If, for example, there are two bidders of the same type, $\bar{b} = \mathbb{E}[v]$. Therefore, modeling the bid distribution as Weibull implicitly imposes an infinite mean on bidder values. We view this problem as largely technical because it results from a very small fraction of large bids being rationalized with implausibly high values. Our solution therefore is to truncate the estimated bid distributions.⁴⁴

To identify maximum bids at which to truncate, we exploit two facts. First, truncating the bid distribution does not affect the reverse hazard rate g_k/G_K , and hence leaves the estimated inverse bid function $\phi(\cdot)$, defined in (11), unchanged for bid values below the truncation. Second, the estimated bid function $\phi^{-1}(\cdot)$ becomes very flat for high bidder values. This means that if we use our prior knowledge of timber auctions to specify a plausible maximum value and use the estimated bid function to locate the implied maximum bid, our resulting truncation point will be relatively insensitive to the precise maximum value we specify.

To make this operational, we observe that values in our model take the form: $v_{it} = \exp(X_t \beta_X + N_t \beta_N) \cdot \xi_{it}$. We assume that for the “stronger” bidder type in a given auction (i.e. mills if any are present, otherwise loggers) $\mathbb{E}_{X_t}[\exp(X_t \beta_X + N_t \beta_N)] \cdot \xi_{it} \leq 3000$, so that for the average tract in our sample, the highest possible value is \$3000 per

⁴⁴An alternative would be to specify directly a bid distribution with finite support, but this has serious pitfalls 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.

mbf. This assumption implies an upper bound on the value distribution $\bar{v}_t(X_t, u_t, N_t)$:

$$\bar{v}(X_t, u_t, N_t) = 3000 \cdot \frac{\exp(X_t\beta_X + N_t\beta_N)}{\mathbb{E}_{X_t}[\exp(X_t\beta_X + N_t\beta_N)]}.$$

For an auction with a set n_t of participants, the bid resulting from this maximum value, $\bar{b}(X_t, u_t, N_t, n_t)$, satisfies:

$$\phi_M(\bar{b}(X_t, u_t, N_t, n_t); X_t, u_t, N_t, n_t) = \bar{v}_k(X_t, u_t, N_t).$$

We calculate $\bar{b}(\cdot)$ numerically for each (X_t, u_t, N_t, n_t) and truncate the bid distribution. If both mills and loggers participate, this truncation also impose an upper bound on logger values, one that may be below $\bar{v}(\cdot)$. In practice, we end up truncating only a very small fraction of the bid distribution. In the auction plotted in Figure 3, for instance, less than 1% of mill bids and 0.001% of logger bids are truncated.

A slight concern with our procedure is that the truncation is imposed *after* we estimate the bid distribution. One way to view what we do is as the first step of an iterative process where we repeatedly estimate the bid distributions, calculate $\bar{b}(X, u, N, n)$, and then re-estimate the bid distributions imposing the new truncation. Because our one-step procedure leads us to truncate such a small fraction of bids, we believe that iterating the procedure would lead to extremely similar estimates.

References

- Abadie, A. and G. Imbens**, “Large Sample Properties of Matching Estimators for Average Treatment Effects,” UC-Berkeley Working Paper, 2004.
- Athey, S. and P. Haile**, “Identification in Standard Auction Models,” *Econometrica*, November 2002, 70 (6), pp. 2107-2140.
- 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.
- Arozamena, L. and E. Cantillon**, “Investment Incentives in Procurement Auctions,” *Review of Economic Studies*, January 2004, 71(1), pp. 1-18.
- 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.

