

NBER WORKING PAPER SERIES

COMPARING OPEN AND SEALED BID AUCTIONS:
EVIDENCE FROM TIMBER AUCTIONS

Susan Athey
Jonathan Levin
Enrique Seira

Working Paper 14590
<http://www.nber.org/papers/w14590>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2008

We thank Phil Haile, Jerry Hausman, Ali Hortacsu, Guido Imbens, Richard Levin, Paul Milgrom, Ilya Segal, a co-editor and our referees for helpful suggestions. We are especially grateful to Rob Porter for detailed and insightful comments. Dominic Coey provided outstanding research assistance. Athey and Levin acknowledge the support of the National Science Foundation and Levin thanks the Alfred P. Sloan Foundation. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Susan Athey, Jonathan Levin, and Enrique Seira. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions
Susan Athey, Jonathan Levin, and Enrique Seira
NBER Working Paper No. 14590
December 2008
JEL No. D44,D82,L13,L40

ABSTRACT

We study entry and bidding patterns in sealed bid and open auctions with heterogeneous bidders. Using data from U.S. Forest Service timber auctions, we document a set of systematic effects of auction format: sealed bid auctions attract more small bidders, shift the allocation towards these bidders, and can also generate higher revenue. We show that a private value auction model with endogenous participation can account for these qualitative effects of auction format. We estimate the model's parameters and show that it 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.

Susan Athey
Department of Economics
Harvard University
1875 Cambridge St.
Cambridge, MA 02138
and NBER
athey@fas.harvard.edu

Enrique Seira
Government of Mexico
Alfonso Reyes No. 30, Piso 12
Col. Condesa, C.P. 06140
México, D.F.
eseira@economia.gob.mx

Jonathan Levin
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
jdlevin@stanford.edu

1. Introduction

The relative performance of open and sealed bid auctions is a central issue in auction design. The choice between these auction formats arises frequently: in allocating natural resources, in procurement, in sales of art, real estate and other assets. 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. But in practice the assumptions of competitive risk neutral bidders with independent identically distributed values often seem too strong. Further theoretical work shows that auction choice becomes relevant as these assumptions are relaxed. The comparison between open and sealed bidding then depends on both the details of the market (e.g. bidder heterogeneity, entry costs, collusion, correlation in bidder values, risk-aversion, transaction costs) and the designer’s objective (e.g. revenue maximization or efficiency).

This wealth of theory cries out for empirical evidence, but it has arrived slowly. 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 provide some new evidence from sales of timber in the national forests. The U.S. Forest Service timber program provides an excellent test case in market design as it uses both open and sealed bidding, 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.

We analyze data from open and sealed bid sales held during 1982 and 1990 in two areas: the Idaho-Montana border and California. We document significant departures from revenue equivalence in terms of bidder participation, allocation and prices. Conditional on sale characteristics, sealed bid auctions induce more participation by small firms that lack manufacturing capacity (“loggers”). In contrast, entry by larger firms with manufacturing capability (“mills”) is roughly the same across auction formats. Sealed bid auctions also are more likely to be won by loggers. Finally, we measure winning bids to be about 10% higher in the sealed bid auctions in the Northern forests. In the California forests, the price difference is small and statistically insignificant.

Motivated by these findings, we consider a model that incorporates several salient departures from the standard independent private value auction model (the revenue equivalence

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).

benchmark). First, we allow bidders to have heterogeneous value distributions. Second, we explicitly model participation by assuming it is costly to acquire information and bid in the auction. Third, because of the large price differential in the Northern forests, we entertain the possibility that mills behave cooperatively in the open auctions. Collusion has been a long-standing concern in timber auctions and the prevailing view is that open auctions are more prone to bidder cooperation because participants are face-to-face and can react immediately to opponents' behavior.

The theoretical effect of bidder heterogeneity is well-known from the work of Maskin and Riley (2000). In an open auction, the bidder with the highest value always wins. But in a sealed bid auction, relatively strong bidders have greater incentive to shade their bids below their true valuations, so a weak bidder can win despite not having the highest valuation. The resulting distortion tips sealed bid outcomes toward weaker bidders and provides them with an extra entry incentive. Using our bidding data, we provide sharp evidence that mills are systematically stronger bidders, so the basic predictions of the heterogeneous bidding model concerning allocation and entry go in the same direction as our empirical findings. The theory is less clear-cut on the relative prices in open and sealed auctions. As Maskin and Riley observed, the comparison hinges on the model primitives: the bidders' value distributions and the cost of participation.

To assess whether the model can match our empirical findings about auction prices, as well as our quantitative findings about allocation and entry, we estimate the structural parameters of the model using data from the sealed bid auctions. We use a parametric version of Guerre, Perrigne and Vuong (2000) to recover the distributions of bidder values from the observed bids. We also recover estimates of entry costs by estimating the distribution of logger entry and combining this with the post-entry profits implied by the estimated value distributions. We show informally and using formal specification testing that the model provides a good fit to the sealed bid data. The estimates indicate substantial differences between mills and loggers, as well as fairly low entry costs and profit margins.

We then use the calibrated model to predict the outcomes of the open auctions in our data under alternative behavioral assumptions. The predictions are out-of-sample in two directions: we predict outcomes for different sales and for a different auction game than was used in estimation. Nevertheless when we compare the model's predictions to the actual auction outcomes, we find that the model plausibly explains the observed differences in participation and allocation across auction formats. Our baseline assumption of competitive bidding also fits the auction prices in California quite well. The competitive benchmark has

a harder time explaining the large price difference between open and sealed bidding that we observe in the Northern forests. Instead the data appear consistent with a mild degree of cooperative behavior by participating mills.

The calibrated model also permits a welfare assessment of the choice between open and sealed bid auctions. We find that for a fixed set of participants, the model predicts relatively small discrepancies in outcomes. 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: sealed bidding increases revenue by roughly 2-5% relative to a competitive open auction due to increased logger entry. Strikingly, even a mild degree of cooperative bidding 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 5-10%). This suggests that bidder competitiveness merits considerable attention in the choice of auction format.

Our paper contributes generally to the economics of auction design, and to several more focused literatures. In particular, a long-standing debate surrounds the format of federal timber sales. Mead (1966) argued early on that open timber auctions generated less revenue. In 1976, forests in the Pacific Northwest, which has historically used open auctions, ran a number of sealed bid sales. Johnson (1979) and Hansen (1986) studied this episode and reached conflicting conclusions. Johnson finds higher prices in the sealed bid auctions, while Hansen argues that the differences are insignificant after accounting for sale characteristics. As Hansen points out, however, the episode is not an ideal testing ground. The choice of auction format was sensitive to lobbying, creating a potentially severe endogeneity problem, and moreover, one might be skeptical of drawing conclusions from 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. Our paper expands on this prior work by addressing a broader set of questions about allocation and participation as well as prices, and in estimating a tightly-specified theoretical model, but we have drawn on Schuster and Nicolucci's work in constructing our sample.

Our work also relates to the empirical literature on bidder collusion. Researchers have proposed several approaches to assess whether auction data are consistent with competitive or collusive bidding (Porter and Zona 1993, 1999; Baldwin, Marshall, and Richard, 1997;

Bajari, 1997; Pesendorfer, 2000; Bajari and Ye, 2003; Asker, 2008). These approaches either require prior knowledge about the existence and structure of a cartel, or derive in-sample specification tests of the competitive model and treat collusion as the alternative. Our method differs in that we use behavior in one auction format as a benchmark from which to evaluate the competitiveness of behavior under an alternative format.

Finally, our empirical approach to studying bidder participation shares features with the industrial organization literature on entry and market structure (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. Bajari and Hortacsu (2003), Li (2005), Li and Zheng (2006), and Krasnokutskaya and Seim (2005) are other recent auction studies that account for bidder participation.

2. The Theoretical Model

This section develops the theoretical model we use to frame our empirical analysis. Our starting point is the heterogeneous private values auction model of Maskin and Riley (2000), which we slightly extend to incorporate participation decisions and possible collusion in open auctions. We discuss some specific modeling choices at the end of the section.

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 must incur a cost K to gather information and enter the auction. By paying K , bidder i learns his (private) value for the tract, v_i , and may bid in the auction. We refer to bidders who acquire information as *participants*, and denote the set of participants by n .

We assume each bidder i 's value is an independent draw from a distribution F_i with continuous density f_i and support $[\underline{v} = r, \bar{v}_i]$. 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 that for all v ,

$f_M(v)/F_M(v) \geq f_L(v)/F_L(v)$. We sometimes refer to the mills as strong bidders and the loggers as weak bidders.

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.

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 p_i specifies a probability of entering the auction.

A *type-symmetric entry equilibrium* is a pair of bidding strategies $b_L(\cdot; n), b_M(\cdot; n)$ and entry strategies p_L, p_M with the property that: (i) loggers use the strategy b_L, p_L and mills the strategy b_M, p_M ; (ii) each bidder's bid strategy maximizes his profits conditional on entering; and (iii) each bidder enters if and only if his expected profit from entry exceeds the entry cost (and may enter probabilistically if the two are equal). To characterize type-symmetric equilibria, we first consider the bidding game and then the entry game.

B. Equilibrium Bidding

We begin with the sealed bid auction. Suppose i is a participating bidder with value v_i , and the set of participants is n . Bidder i 's 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 $G_j(b; n) = F_j(b_j^{-1}(b; n))$ is the probability that j will bid less than b . 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 for any set of participants n , and provide a basis for estimating bidders' value distributions (Guerre, Perrigne and Vuong, 2000).

Equilibrium bidding behavior exhibits several key features. First, bid strategies are type-

symmetric. Second, mills submit higher bids: $G_M(b; n) \leq G_L(b; n)$ for all b . This is a testable implication of the model. Third, mills shade their bids more than loggers: $b_M(v; n) \leq b_L(v; n)$ for all v . This implies that a logger may win despite not having the highest value.

Now consider the open auction. In this case, all participants have a dominant strategy of bidding up to their valuation, so $b_i(v; n) = v$. Bidder i 's expected profit conditional on entering and having value v_i is

$$\pi_i^o(v_i; n) = \max_{b \geq r} (v_i - \mathbb{E}[\max\{v_{-i}, r\} | v_j \leq b \ \forall j \in n \setminus i]) \prod_{j \in n \setminus i} F_j(b). \quad (3)$$

Unlike the sealed bid auction, the open auction is efficient: the entrant with the highest value wins the auction.

C. Equilibrium Entry

We now characterize equilibrium entry. Let $\pi_L^\tau(n_L, n_M)$ and $\pi_M^\tau(n_L, n_M)$ denote the expected logger and mill profit in auction format $\tau \in \{o, s\}$ if the set of participants n includes n_L loggers and n_M mills, and participants use equilibrium bid strategies. Then bidder i 's ex ante expected profit from participating is

$$\Pi_i^\tau(p) = \sum_{n \subset N} \pi_i^\tau(n_L, n_M) \Pr[n_L, n_M \mid i \text{ enters, opponents play } p_{-i}], \quad (4)$$

where $p = (p_1, \dots, p_{M+L})$ is the profile of entry probabilities, and π_i^τ equals π_L^τ or π_M^τ depending on whether i is a logger or mill. Entering is optimal if the expected profit $\Pi_i^\tau(p)$ exceeds the entry cost K .

A type-symmetric entry equilibrium (p_L, p_M) exists for both auction formats, but in general it need not be unique. The following result is useful in this regard.

Proposition 1 *Suppose that for all n_L, n_M , $\pi_M^s(n_L, n_M + 1) > \pi_L^s(n_L, n_M)$. Then there is a unique type-symmetric entry equilibrium for both auction formats. In equilibrium, either $p_L = 0$ or $p_M = 1$.*

The uniqueness condition requires that mills have a sufficient value advantage over loggers to outweigh the effects of facing an additional bidder. As a matter of theory it is rather strong. In our empirical work, however, we estimate bidder value distributions without making any equilibrium assumptions about entry behavior, and then verify that the condition holds for each sale tract in our data. Thus the calibrated version of our model always has a unique

type-symmetric entry equilibrium. In our data, we observe logger entry in more than 85% of sales and always more potential logger entrants than actual logger entrants, so the empirically relevant equilibrium appears to be one in which each logger enters with probability between zero and one.²

D. Comparing Auction Formats

We now compare the equilibrium outcomes of open and sealed bid auctions. A useful benchmark to have in mind is the case where bidders are homogeneous, so $F_L = F_M$. In this case, we have auction equivalence as follows: *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, (ii) the same allocation; and (iii) the same expected revenue.*

This equivalence breaks down with heterogeneous bidders. Because mills shade their bids more than loggers in the sealed bid equilibrium, a logger has a greater chance to win, and hence has greater expected profits, than in an open auction where the allocation is efficient. The argument is reversed for mills, leading to the following result.

Proposition 2 *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.*

The statement of the result is complicated by the possibility of multiple equilibria. Under the uniqueness condition of Proposition 1, however, the prediction is unambiguous: open bidding leads to less logger entry, equivalent mill entry and a lower chance that a logger wins.³

There is no general theoretical comparison of expected revenue, even with fixed participation. This provides further motivation for the parameterized model we develop in Section 5. The model does imply that sealed bidding is less efficient. The sealed bid auction is inefficient even conditional on participation, while the socially efficient type-symmetric outcome is achieved as an equilibrium of the open auction (Athey, Levin and Seira, 2004).

²The condition in Proposition 1 also greatly restricts the set of non-type-symmetric equilibria as it implies that in *any* equilibrium where *any* logger enters with positive probability, *every* mill must enter with probability one. If for example, we were to restrict attention to pure strategy entry equilibrium, every equilibrium would involve some number n_M of entering mills and n_L of entering loggers, where $n_L > 0$ would imply $n_M = N_M$.

³The idea that sealed bidding may increase entry by weaker bidders is emphasized by Klemperer (2004) in the context of spectrum auctions.

E. Collusion in Open Auctions

As noted in the Introduction, we estimate that in some forests open auction prices are substantially lower than sealed bid prices. This finding, and the fact that collusion in open auctions has been a long-standing concern in Forest Service sales (Mead, 1966; U.S. Congress, 1976; Froeb and McAfee, 1988; Baldwin et al, 1997), suggests incorporating open auction collusion into the model.

Collusive schemes can take many forms, so we assume for concreteness that participating mills at an open auction cooperate perfectly. 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 cooperating with other participating mills, but do not coordinate entry.⁴

Fixing the set of participants, collusion clearly will lower prices and increase mill profits. But 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 may crowd out logger participation.

Proposition 3 *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.*

For the empirically relevant case in which there is a unique type-symmetric equilibrium where mills enter with probability one and loggers randomize, collusion has no effect on entry or allocation relative to the competitive open auction outcome. It simply lowers prices. Therefore to the extent that the competitive model might explain observed departures from revenue equivalence in terms of entry and allocation, the possibility of collusion provides further flexibility in terms of explaining price differences across auction formats.

F. Discussion of Modeling Choices

Our model omits at least two forces of potential importance: common values and bidder risk-aversion. In timber auctions, differences in bidder costs and contractual arrangements

⁴There 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.

provide a source of private value differences. At the same time, bidders can obtain private estimates of the quality and quantity of timber and may have differing beliefs about future market conditions. This suggests a potential “common value” component as well (Athey and Levin, 2001; Haile, Hong and Shum, 2003).⁵ It is also plausible that bidders at Forest Service timber auctions exhibit a degree of risk-aversion. Indeed Athey and Levin (2001) provide some indirect support for this based on the way observed bids are constructed (see also Perrigne, 2003).

While a model that allows for common values and bidder risk-aversion might have additional realism, we decided to abstract away from them for two reasons. First, as we discuss below, our empirical results indicate departures from revenue equivalence that are qualitatively different from those implied by common values or risk-aversion. Second, incorporating either significantly complicates the analysis. Hence we opted to use a simpler model we felt might still explain the data.

3. Timber Sales

We now describe the key institutional features of timber auctions, our data, and the process through which the Forest Service decides when to use open or sealed bidding.

A. The Timber Sale Process

Our data consists of timber sales held between 1982 and 1990 in Kootenai 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. We also provide evidence from sales held in California between 1982 and 1989. These 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

⁵Athey 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. Haile (2001) analyzes how resale markets can lead to common values even if the underlying environment is one of private values.

publicly at least thirty days prior to the auction. The announcement includes the form of the auction, estimates of available timber and logging costs, tract characteristics and a reserve price. The 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. In some cases, the Forest Service restricts entry to firms with less than 500, or less than 25, employees. We do not consider these small business 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. 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. A useful consequence is that we observe all open auction bidders, even those who do not bid actively, which allows a comparison of entry patterns across auctions.

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. The bidders in these auctions range from large vertically integrated forest products conglomerates to individually-owned logging companies. To study participation and allocation in a way that respects this variation, we classify bidders into two groups: “mills” that have manufacturing capacity and “loggers” that do not. One can imagine other possible ways to try to capture the diversity of bidders. In practice, however, other natural groupings, such as by number of employees or by number of auctions entered, turn out to be quite similar.

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 regress the per-unit bids (in logs) from the sealed bid auctions on a

dummy for whether the bidder is a mill and auction fixed effects. For the Northern forests, we estimate a mill dummy coefficient of 0.248, meaning mill bids are 25% higher on average, with a t -statistic of roughly 8. An entering mill is also more likely to win than an entering logger (28% versus 21%). The pattern in California is similar though the magnitudes are smaller. Controlling for auction fixed effects, mill bids are just over 12% higher on average. Mills are also more likely to win conditional on participating in an auction.

The model of entry requires that we have a measure of potential bidders for each sale. As we discuss below, so long as loggers enter with positive probability (the relevant case for our data), our calibrated model implies that in any equilibrium all mills must enter. Therefore under the assumption that firms use equilibrium entry strategies, we can infer that the number of potential mill entrants for a given sale is just the number of actual mill entrants. To construct a measure of potential logger entrants for a given sale, we count the number of distinct logging companies that entered an auction in the same geographic area in the prior year.⁶ We also do a similar count for mills and use it as a control in our baseline regressions, where we do not impose any assumption of equilibrium behavior.

Table 1 presents summary statistics of sale characteristics and auction outcomes. Focusing on the full sample, there are some obvious differences between the open and sealed bid auctions. In the Northern forests, the average sale price per unit of timber (in 1982 dollars per thousand board feet of timber or \$/mbf) is roughly \$62 in the sealed auctions and \$69 in the open auctions. The number of entering logging companies is also somewhat higher in sealed auctions (3.2 versus 2.5), 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. Sale differences, particularly in terms of size are even more pronounced in California. 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

⁶This measure probably suffers from a degree of measurement error. Firms may go in and out of business or become more or less active in Forest Service auctions over time without our knowledge. Moreover, the Forest Service data records bidder names with a variety of spellings and abbreviations. Despite carefully checking each name and cross-referencing with industry reference books, we may not have obtained perfectly accurate counts. Note that for the large Northern forests, we use forest-district as the relevant geographic area; for California we use forest.

The U.S. Forest Service has historically used both open and sealed bid auctions to sell timber from the national forests. The policy in place at the time of our data arose from a debate that followed passage of federal legislation in 1976. At that time, Congress proposed the use of sealed bidding. The implementation of the law, however, allowed individual forest managers to use open auctions if they could justify the choice. As a result, sale method has varied geographically. In the Pacific Northwest, for instance, open auctions predominate. We focus on areas that have used a more balanced mix of open and sealed bidding.

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 have replicated our analysis using the subsample that Schuster and Niccolucci (1993) identify as randomized, and get similar results to what we report below, though our estimates lose precision due to the smaller sample size.⁷

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, the U.S. Census count of the number of logging firms and sawmills in the county of the sale, and our measure of potential bidders. In addition, we include dummy variables for the year of the sale, the quarter of the sale, the area in which the sale took place (forest district in the Northern region and forest in California), and if major species were present. We are particularly sensitive to the importance of sale size, so rather than simply assuming a linear or quadratic effect, we specify its effect as a step function with 10 steps that roughly correspond to deciles in the data.

⁷Relative to Schuster and Nicolluci, we use more districts and years within the two largest Northern Region forests (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. Schuster and Nicolluci, however, include some sales from other forests. We focus on the two largest 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.

The results are reported in Table 2. As expected, sale size is a significant correlate of auction method, particularly in California. 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 forests and 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 the predicted probability of being sealed according to our logit regression, i.e. the propensity score, will be close to zero or one. This occurs for many of the open auctions in California, mainly because in that region very large sales are almost certain not to be sealed bid.

In order to compare relatively similar tracts in our empirical analysis, we drop sales that have a propensity score below 0.075 or above 0.925. This results in dropping 154 open auctions and 8 sealed auctions in the Northern Forests. It has a much more dramatic effect in California, where we retain only one-third of the sales. The result, however, is that the selected sample has much smaller differences in sale characteristics across sale format.

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

$$Y = f(SEALED, X, N, \varepsilon), \tag{5}$$

where $SEALED$ 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 represents measures of potential competition, and ε is unobservable. We are interested in the average effect of auction format, denoted $\tau_Y = \mathbb{E}_{X,N,\varepsilon}[f(1, X, N, \varepsilon) - f(0, X, N, \varepsilon)]$.

The crucial identifying assumption auction format is independent of the unobserved component ε conditional on covariates. This clearly holds for sales where auction method was randomly designated, although it is important that X included the administrative unit doing the randomization, given that assignment probabilities differed by forest district. 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.⁸

We consider three alternative estimates of the “average treatment effects” τ_Y . The first is an ordinary least squares regression:

$$Y = \alpha \cdot SEALED + X\beta + N\gamma + \varepsilon, \quad (6)$$

which is easily interpretable but doesn’t allow the effect of sealed bidding to vary across tracts. The second specification allows for this variation by interacting *SEALED* with the individual covariates. We then compute and report an average effect for the sample. The third approach is a matching estimator that matches every sealed bid auction with the M “closest” open auctions and vice versa. Closeness is measured by distance between the estimated propensity scores of the auctions in the sample.⁹ The average effect of auction format is calculated by comparing the outcome of each sale t , Y_t , with the average outcome the matched sales \hat{Y}_t :

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

Here T is the number of sales. We implement this estimator, setting $M = 4$, and compute robust standard errors following Abadie and Imbens (2006). The three alternative approaches yield very similar empirical results, providing assurance that our findings are not driven by

⁸If the forest manager uses a deterministic rule, for instance using open auction whenever 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.

⁹We also experimented with using larger numbers of sale characteristics in constructing matches, and with adjusting for bias as suggested by Abadie and Imbens. To do this we define the distance between sales with covariates x and z as $\|x - z\|_W$, where $\|x\|_W = (x'Wx)^{1/2}$ and W is a diagonal matrix consisting of the inverses of the variances of the covariates x . There is some sensitivity to the exact choice of matching covariates and use of bias correction, and alternative matching strategies arguably suggest larger effects of sealed bidding than our reported estimates. We report the propensity score match estimates as they are conservative and in line with the regression estimates.

a particular specification or functional form assumption.

B. Evidence from the Northern and California Forests

We report our empirical results on the effect of auction choice in Table 3. Each column displays the estimated effect of sealed bidding on a sale outcome conditional on sale characteristics, with the relevant outcomes being logger entry, mill entry, bidder composition and sale revenue.

We find that sealed bidding has a strongly positive effect on logger entry in both the Northern and California forests. In particular, we estimate that sealed bid auctions attract around 10% more logger entrants in both the Northern and California forests. This translates into roughly 2-3 additional entrants for every 10 sales. All six point estimates are statistically significant; the estimates are somewhat more precise in the Northern forests where the sample is larger. In contrast, sale format appears to have little effect on entry by mills. All specifications for the Northern forests, and the regression specifications for California yield small and statistically insignificant effects. The one exception is the matching estimate for California, which suggests lower mill participation in the sealed bid auctions.¹⁰

The consequence of increased logger participation and unchanged or decreased mill participation is that the composition of bidders in sealed bid auctions is shifted toward loggers. We estimate that the fraction of participants who are loggers 5-6% higher in sealed bid auctions in both the Northern and California forests. The composition effect suggests that sealed bid auctions will be more likely to be won by loggers. Our findings are consistent with this as well. We estimate a 3-4% greater chance that a logger will win if the auction is sealed bid. These last point estimates are not highly precise, particularly in California, so we cannot rule out a fairly small effect of auction format on allocation.

The final columns of Table 3 report our estimates of the effect of auction format on the sale price per unit volume. Here our findings differ dramatically across the two areas. In California, we find little difference in sale price between the two auction formats. Our estimates indicate slightly higher revenue in the sealed bid auctions, but the finding is not statistically significant and reverses after controlling for the number of entering loggers and mills. In the Northern forests, however, we find that sealed bid prices are around 10% higher than open auction prices after controlling for sale characteristics. Our point estimates are highly significant. To get a sense of the magnitude of this effect in dollar terms, note that the

¹⁰Although we will not develop the point, we note that reduced mill participation in sealed bid auctions would be consistent with a version of the theoretical model where entry costs are heterogeneous (Athey, Levin and Seira, 2004).

average winning bid (in 1982 dollars rather than 1982 dollars per unit volume) is just over \$134,000. So a 10% difference in the winning bid price translates into a \$13,000 difference in Forest Service revenue per sale, or about \$14 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 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 6% (s.e. 3%) more revenue. The matching estimate is a bit higher at 9%. 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.¹¹

C. Explaining the Departures from Revenue Equivalence

At a qualitative level, the theoretical model developed earlier is consistent with all of the empirical findings just reported: greater logger participation in sealed bid auctions, a negligible change in mill participation, a higher probability in sealed bid sales that a logging company will win, and either a small difference in prices across auction formats or substantially higher prices in the sealed bid sales. Moreover, the key assumptions generating these departures from revenue equivalence: that bidders are heterogeneous, that mills are stronger bidders than logging companies and that entry should be treated as endogenous, also seem consistent with the data.

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. With this motivation, we turn in the next section to estimating the model's parameters and comparing the quantitative predictions of the theories to the data.

Before doing this, however, we pause to consider whether there might be alternative

¹¹A natural concern in interpreting this revenue decomposition arises if there are sale characteristics that are observed by the bidders prior to making their entry decision but not accounted for in our data. In this event, the number of entrants is endogenous in this regression. To explore this, we experimented with using our measures of potential competition as an instrument for the number of entering bidders. We found, however, that our estimated coefficients were highly sensitive to the particular choice of potential competition measures, none of which are ideal.

explanations for our empirical findings that are distinct from the forces captured in our theoretical model. One possibility is that our estimates do not reflect the systematic effects of auction format at all, but rather a 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, but contradicts both perceptions within the industry and the Forest Service's own guidelines. Industry lore suggests a scenario where the mills prefer oral auctions (as predicted by our theory), and where forest managers defer to the mill's preferences. And the Forest Service instructs managers to use sealed bidding if they expect a sale *not* to be competitive (Forest Service Handbook 2409.18, Chapter 57.1).

Another possibility is that our finding do reflect systematic departures from auction equivalence, but not for the reasons captured in our model. For instance, our 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 their presence, neither seem likely to be the primary source of the departures we observe from revenue equivalence. Theoretical models with common or affiliated values (and without the other elements of our model, namely bidder heterogeneity and collusion), imply lower prices in sealed auctions rather than higher as we observe in the data. Bidder risk-aversion potentially could explain the observed prices, But at least in the cases considered by Matthews (1987), risk aversion would lead to lower participation in sealed bid auctions contrary to our findings. So to the extent that either common values or bidder risk-aversion would help to explain the data, they would have to be part of a more complicated story.

5. Structural Estimation and Testing

In this section we bring the model and the data together to assess the relationship between our empirical 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 the tract characteristics. To do this, we assume competitive behavior in the sealed bid auctions as outlined above. We allow for both observed and unobserved heterogeneity in 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 auction tracts were not used to estimate the parameters of the model and because the open auction is a different game than the sealed bid auction around which estimation is based.¹² 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.

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.

A notable feature of our data is that bids within a given sealed bid auction are highly correlated conditional on observed sale characteristics. We therefore follow Krasnokutskaya

¹²In principle one might try to use the data from the open auctions to help estimation the model. Athey and Haile (2002), however, show that when values are correlated as in our model of unobserved heterogeneity, underlying value distributions cannot be identified from open auction bids. Haile and Tamer (2003) point out additional concerns with drawing inferences from losing bids in open auctions.

(2004) in allowing for unobserved heterogeneity in sale characteristics. An extension along these lines appears crucial as, in line with Krasnokutskaya’s work on highway procurement, we estimate implausibly high bid margins when we fail to account for within-auction bid correlation.¹³

Formally, let X denote the set of sale characteristics known both to the econometrician and the bidders. Let u denote an auction characteristic known to participating bidders but not observed in our data. Let $N = (N_L, N_M)$ represent the number of potential mill and logger entrants. And let $n = (n_L, n_M)$ denote the numbers of participating mills and loggers. We assume that bidders initially have the information in the sale announcement and knowledge of the set of potential bidders; that is, they know (X, N) . They then decide whether to incur the entry cost, $K(X, N)$, and participate in the auction. If they participate, they learn the set of participating bidders n , the sale characteristic u and their private value. We write bidder value distributions as $F_L(\cdot|X, u, N)$ and $F_M(\cdot|X, u, N)$, and assume that values are independent conditional on (X, u, N) .

Given these assumptions, we can write the equilibrium bid distributions as $G_L(\cdot|X, u, N, n)$ and $G_M(\cdot|X, u, N, n)$. We assume that if there is a single bidder he optimally bids the reserve price, but otherwise we treat the reserve price as non-binding.¹⁴ More generally, we assume the data we observe is generated by a type-symmetric entry equilibrium. As we discuss below, there is a unique such equilibrium consistent with the estimated value distributions and observed entry probabilities. In this equilibrium, mills enter with certainty and each logger enters with some probability between zero and one depending on sale characteristics. This means that we can infer the number of potential mill entrants N_M as equal to the number of participating mills n_M . For each sale, we use our count of active logging companies described earlier as our measure of potential logging entrants, N_L . 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

¹³An alternative way to rationalize correlation in bids is with an affiliated private values model, but at least in the baseline symmetric model affiliation implies that prices will be higher in open auctions, contrary to our data. As an institutional matter, we also believe it plausible that bidders commonly observe certain features of a tract that make it more or less valuable.

¹⁴See 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.

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 correlation of the bids. We adopt a parametric approach to estimate these three distributions.

Our particular model specifies Weibull bid distributions with Gamma distributed auction heterogeneity. Thus we assume that for $k = L, M$:

$$G_k(b|X, u, N, n) = 1 - \exp\left(-u \cdot \left(\frac{b}{\lambda_k(X, N, n)}\right)^{\rho_k(n)}\right). \quad (7)$$

Here $\lambda_k(\cdot)$ is the scale, and $\rho_k(\cdot)$ 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 $\ln \rho_k(n) = n\gamma_{n,k} + \gamma_{0,k}$.^{15,16} We assume u has a Gamma distribution with unit mean and variance θ , and is independent of X, N , and n . We estimate these parameters of the model, (β, γ, θ) , by maximum likelihood; the likelihood function is written out in the Appendix. The estimates are reported in Table 4.

Several points about the estimated bid distributions deserve mention. First, recall that the basic assumption of the theory was that mill values stochastically dominate logger values, and an implication was that mill bids should dominate logger bids. Our empirical specification does not impose this. Nonetheless, we find that mill bids do dominate those of loggers. On average, mill bids are roughly 25% higher than logger bids in the Northern forests and 15% higher in California. 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 for both geographical regions that u has significant variance, confirming that our modeling of unobserved heterogeneity across auctions is warranted.

Importantly, the Gamma-Weibull functional form appears to provide a good fit to the observed distribution of logger and mill bids, the within-auction bid correlation, and the observed sealed bid prices. Our model has the useful property bidder i 's bid in auction t

¹⁵The 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.

¹⁶Specifying 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 $\bar{n} = 5$.

can be expressed as $b_{it} = \exp(X_t\beta_X + N_t\beta_N) \cdot \varepsilon_{it}(n)$. Defining the sealed bid residuals as $\hat{\varepsilon}_{it} = b_{it} / \exp(X_t\hat{\beta}_X + N_t\hat{\beta}_N)$, we investigate how closely these residuals match the distribution of the ε_{it} 's predicted by our fitted model. In the Northern forests, the overall mean of the bid residuals is 2.16; the standard deviation is 1.18; the between-auction standard deviation is 0.94 and the within-auction standard deviation is 0.75. By way of comparison, the fitted model predicts a mean of 2.12, and respective standard deviations of 1.21, 0.97 and 0.70. We obtain a similarly close fit in the California forests, where the respective numbers from the data are 25.9, 12.6, 9.7 and 8.5 compared to our model's prediction of 25.6, 13.6, 10.8 and 8.1. To provide a visual picture, Figure 1 plots the distribution of sealed bid residuals in our sample (i.e. the distribution of the $\hat{\varepsilon}_{it}$ s, where) next to the distribution predicted by our fitted model.

Despite this informal confirmation of model fit, one might still wonder whether our parametric modeling is unduly restrictive.¹⁷ To address the issue more formally, we implement a natural specification test due to Andrews (1997). Andrews' Conditional Kolmogorov Test tests the null hypothesis that conditional on a set of exogenous covariates, a set of endogenous variables is generated by a particular parametric distribution. In our case the exogenous covariates are the sale characteristics (X, N) , the endogenous variables are the bids, and the parametric model is the Gamma-Weibull mixture model. Andrews' test is based on a bootstrap procedure in which one uses the estimated model to repeatedly draw samples of the endogenous variables and compares these simulated datasets to the observed data. We implement the test and find that we cannot reject the null hypothesis that our parametric specification is correct, even at very high confidence levels (20% in both the Northern forests and California). These findings provide additional support for our modeling approach.

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

¹⁷For instance, Krasnokutskaya (2004) estimates a semi-parametric model with unobserved heterogeneity assuming that the unobserved component of the bid separates multiplicatively into an auction effect and an idiosyncratic component. In our setting, an important practical problem with semi-parametric estimation is that one would want to estimate the model separately for each vector of participants (n_L, n_M) and we simply don't have the data to do this. Papers that use a parametric strategy include Jofre-Bonet and Pesendorfer (2003) and Krasnokutskaya and Seim (2006). The latter follow our lead in using a parametric model with unobserved heterogeneity.

b_i through his first-order condition for optimal bidding:

$$v_i = \phi_i(b_i; X, u, N, n) = b_i + \frac{1}{\sum_{j \in n \setminus i} \frac{g_j(b_i|X, u, N, n)}{G_j(b_i|X, u, N, n)}}. \quad (8)$$

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 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 (2004), 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:¹⁸

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

Figure 2 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. To compute the equilibrium bid functions, we combine the fitted bid distributions $G_L(\cdot|X, u, N, n)$ and $G_M(\cdot|X, u, N, n)$, assuming $X = \bar{X}$, $N = \bar{N}$, $u = 1$ and $n = (2, 2)$ with the first-order condition to find $b_k(v|X, u, N, n) = \phi_k^{-1}(v|X, u, N, n)$. 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.

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 in the Northern forests, the median profit margin across all bids is 9.5%. The corresponding number for California is 10.0%. These margins, which are similar when we look separately at mills and loggers, suggest that the sealed bidding is quite competitive.

Finally, we can use the estimated value distributions to investigate whether the equilibrium uniqueness condition in Proposition 1 holds for our calibrated model. Our parametric model has the property that the effect of observed sale characteristics (X, N) is multiplicatively separable. This property that extends from the bids to the bidder values and hence the bidder profits. So we can compute expected equilibrium profits for loggers and mills for

¹⁸A 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.

just a single set of sale characteristics (X, N) and simply re-scale to account for changes in these characteristics.

To compute expected equilibrium profits, we repeatedly simulate the outcomes of sealed bid auctions and average bidder profits over the simulations. In a given simulation, we draw a value for the unobserved auction characteristic u , then sample for each bidder from the estimated distributions G_L, G_M and infer the bidder values that correspond to these draws. This leaves us with bidder values and equilibrium bids so we can identify the auction winner and the realized bidder profits. We simulate 5000 auctions for each plausible level of logger and mill participation (up to eight mills and twenty-five loggers) to compute the expected logger and mill profits, $\pi_L^s(X, N, n)$ and $\pi_M^s(X, N, n)$. These estimates have the property that for all n_L, n_M , $\pi_M^s(X, N, n_L, n_M + 1) > \pi_M^s(X, N, n_L, n_M)$ for every tract in our sample. Therefore Proposition 1 implies the calibrated model has a unique type-symmetric equilibrium for every sale tract irrespective of the fixed cost of entry.

Estimating Entry Costs

The remaining parameter of the model is the entry cost, which we recover using the equilibrium conditions for optimal entry behavior. As just explained, our estimated value distributions imply a unique type-symmetric entry equilibrium with the property that if there is logger entry with positive probability, all mills must enter with probability one. We observe loggers entering 85% of sales in the Northern region and 88% in California, so for these sales we can infer that the number of potential mill entrants equals the number of observed mill entrants, i.e. $N_M = n_M$. For the sales with zero logger entrants, we also make this same inference.¹⁹

As described above, we construct a measure of potential logger entry N_L for each sale by counting the number of loggers entering sales in the same area over the prior year. This number strictly exceeds the number of observed logger entrants in virtually all the sales (99% in the Northern region and 95% in California), indicating that the equilibrium needed to rationalize the data is one in which loggers enter with probability strictly between zero and one. In such an equilibrium, loggers must be just indifferent between entering and not entering. Letting $\Pi_L^\tau(X, N)$ denote the equilibrium profit a logger expects from entering as a function of observed sale characteristics (X, N) and the sale method $\tau \in \{o, s\}$, we have:

¹⁹It is possible that in some of these sales, the relevant equilibrium is one in which the loggers entered with probability zero and perhaps not all mills entered. We assume this is not the case, and perform a specification test, explained below, to test the assumption.

$$\Pi_L^\tau(X, N) = \sum_{n \subset N} \pi_L^\tau(X, N, n) \Pr[n|X, N, i \in n, \tau] = K(X, N). \quad (9)$$

Here $\Pr[n|X, N, i \in n, \tau]$ is the probability that $n = (n_L, n_M)$ bidders enter given that i enters.

Our estimated value distributions already provide an estimate of $\pi_L^\tau(X, N, n)$. We use the sealed bid data to construct an estimate of bidder's beliefs about opponent entry. In equilibrium, $n_M = N_M$, while loggers independently randomize their entry with identical probability $p^s(X, N)$. The distribution of logger entry is therefore binomial, as is the distribution of opponent entry. In particular,

$$\Pr[n_L|X, N, i \in n, s] = p^s(X, N)^{n_L-1} (1 - p^s(X, N))^{N_L-n_L}.$$

For estimation, we specify a parametric model:

$$p^s(X, N) = \frac{\exp(X\alpha_X + N\alpha_N)}{1 + \exp(X\alpha_X + N\alpha_N)}.$$

We estimate the parameter vector α by maximum likelihood using the observed logger entry into sealed bid auctions. These estimates are reported in Table 4.²⁰

Putting the estimated equilibrium profit function $\pi_L^s(X, N, n)$ together with the estimated probability of logger entry $p^s(X, N)$, we use (9) to compute 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 $K(X, N) = \Pi_L^s(X, N)$. We estimate a median entry cost of \$2870 (s.e. \$325) for the Northern forests and \$5056 (s.e. \$673) for the California forests. 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.^{21,22}

²⁰With these estimates in hand, we can check if our assumption of that the probability of logger entry was strictly positive even for the few tracts where we observe zero logger entry. If this were so, we should expect the data to contain significantly more auctions with zero logger entry than is predicted by the binomial model. They do not.

²¹As 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.

²²Our analysis assumes a type-symmetric entry equilibrium. A similar analysis is possible under the assumption that potential entrants play a pure strategy entry equilibrium. In this case, the strong asymmetry between mills and loggers ensures a unique number of mill and logger entrants for any entry cost, and we can use revealed preference to obtain bounds on the fixed entry cost. Proceeding in this fashion, we obtain fairly tight bounds on entry cost for each tract that are similar to the estimates we obtain under the assumption of type-symmetric equilibrium.

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 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 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 the assumption of either competitive and collusive behavior. In practice we do this by 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 procedure gives predicted open auction outcomes for each tract conditional on any hypothetical set of participants. To predict open auction entry, we assume a type-symmetric equilibrium. For each tract we treat mill entry as known and equal to the set of potential mill entrants. We calculate the unique logger entry probability that leaves each logger just indifferent between entering and not entering. This yields the unique equilibrium in logger entry strategies that we combine with our equilibrium bidding predictions to generate predicted outcomes as a function of observed tract characteristics. As was discussed in Section 3, logger entry and auction allocation are the same regardless of whether mill behavior is competitive or collusive; the only difference in outcomes is the predicted auction price.

Table 5 reports the average outcomes in our data and the average outcomes predicted by the parameterized model. We generate standard errors for the predicted outcomes using a parametric bootstrap in which we re-sample from the asymptotic distribution of the bid and entry distribution parameters reported in Table 4 and then repeat the procedure of calculating expected auction outcomes for each bootstrap repetition.

For the Northern forests, the model closely predicts the average auction prices, the average sale revenue and the fraction of sales that loggers win. For instance, the average sale price in the data is \$69.4, while the model predicts an average price of \$70.4, and \$69.9 conditional on the set of participating bidders. The model also predicts the average sealed bids of loggers and mills with reasonable accuracy. The results for the California forests are similarly encouraging. The model closely matches the average logger and mill bids and the fraction of sales won by loggers. Perhaps the biggest discrepancy between the model and the data is that we somewhat overpredict the average sale price and revenue in California relative to the observed outcomes. The average sale price in the data is \$80.4, while the model predicts \$84.4, or \$83.8 if we condition on the participating bidders.

As the model’s parameters are estimated from the sealed bid data, the tight match between predicted and actual outcomes just amplifies our earlier point that the model fits well. The next step, however, provides a demanding test of the theory. We now use the model to predict the outcomes of the open auctions and compare these predictions to the data. Here we are 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 we are predicting the outcomes of a different auction game than was observed in estimating the model’s parameters. These predictions and actual outcomes are reported in the second panel of Table 5.

Strikingly, the model predicts a level of logger entry in open auctions that is very close to the actual level. In the Northern forests, the model predicts an average of 2.67 loggers entering in equilibrium versus 2.75 in reality. In California, the model predicts 1.90 compared to 1.95 in reality. These results indicate that the fitted model can explain the entry differences between open and sealed bid sales in our data that were one of the key departures from revenue equivalence. The model is somewhat less successful in matching the fraction of open auctions won by loggers. In both regions, the model under-predicts how often loggers win. In the Northern forests, for instance, the model predicts loggers will win 54.4% of the open sales, or 56.0% conditional on realized participation, while in reality they win 59.0%. There is a similar discrepancy in California.

Turning to the open auction prices, recall that we observed practically no difference between open and sealed bid prices in California and a sizeable difference in the Northern forests. This observation was part of our motivation for introducing the possibility of open auction collusion into our model. Table 5 shows that for the California forests, the competitive model predicts open auction prices close to the actual prices. The average sale price in the California open auctions was \$85.1. Our fitted model predicts an average price of \$87.2 conditional on realized entry, and \$86.7 when we predict entry as well as bidding. The model therefore seems to replicate our empirical finding of little price differential due to the choice of open or sealed bidding.

The situation is different for the Northern forests where we observed a large price difference between open and sealed auctions. The numbers in Table 5 indicate that observed open auction prices are below the competitive prices predicted by the model, although well above the fully collusive prediction. The competitive model predicts an average price of \$67.8, or \$67.9 conditional on realized entry. The prediction falls to \$44.2 under the assumption that the mills fully collude. In fact, the average sale price across open auctions is \$63.3 per mbf. Accounting for sampling error, we reject both the competitive and collusive models at conventional confidence levels. An assumption of mildly cooperative behavior on the part of participating mills appears to provide a better match than either the competitive or fully collusive extremes.

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. A 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.

As statistical detection of collusion is known to be a difficult problem (e.g. Bajari and Ye, 2003), it is interesting to consider more refined predictions of the collusive model. One such prediction concerns the relationship between prices and the number of participating mills. For sales with zero or one mill, the competitive and collusive model yield identical predictions. Any effect of mill collusion should appear only in sales with more than two mills. To explore this, we divide the sales in the Northern region into three groups: those with zero participating mills, one participating mill, and two or more participating mills. Table 6 then reports the observed and predicted competitive prices for open and sealed sales falling

into these categories. The striking result is that the competitive model predicts prices quite accurately for sales with zero or one mill, but observed open auction prices fall well shy of predicted competitive prices when there are two or more mills. This finding indicates that the price shortfall in Table 5 is driven entirely by sales with multiple mills, consistent with the bidder collusion theory.

C. 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 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 18% of the time (18% being the number that rationalizes the observed open auction prices in the Northern region).

Our comparisons are reported in Table 7, which reports expected auction outcomes taking participation as fixed and solving for the complete entry equilibria under sealed and open bidding. The top panel shows the results for the Northern forests, and the bottom panel for California.

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 \$320 per sale in the Northern region and \$546 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, less than \$100 per sale in both regions.

These differences increase somewhat when we account for the fact that bidder participation will vary systematically with auction format. According to the model, sealed bid and open auctions will attract the same number of mills, but sealed bid auctions will attract between 3-4 more loggers for every 10 sales. 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 2-4% more sales with sealed bidding. A second effect is to increase the revenue advantage of sealed bidding to roughly \$3000 for the average sale in the Northern region and \$14,000 in California. Our estimate of the social surplus differential remains relatively small in the for the Northern region, and is quite noisy for California, to the extent that our point estimate indicates higher social surplus from sealed bidding, despite the fact that we know equilibrium sealed bidding to be less efficient.²³

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 \$14,000 less per sale than competitive sealed bidding if mills are able to engage in a mild amount of cooperative behavior. The difference is over \$17,000 once we account for participation effects. These numbers are even larger on the California tracts. 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.

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. We show that sealed bid

²³The 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 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 has its own nontrivial problems. Notably, for any given set of participants a pooled value distribution estimate does not correspond through the first order condition to the estimated bid distribution. Moreover, because averaging the value distribution estimates leads to a distribution that is flatter than the individual estimates, the resulting sealed bid equilibrium does not match that well with the observed data, which is a main reason why we pursued our current approach.

auctions attract more small bidders, shift the allocation toward these bidders, and in some forests generate higher revenue. We also show that an extension of the standard independent private values auction that can explain these findings, both qualitatively and quantitatively, and furthermore allows us to measure the degree of bidder competitiveness.

Our approach to structural estimation in this setting has two main features. First, motivated by a desire to match key features of the application, we incorporate several elements (heterogeneous bidders, unobserved auction heterogeneity, and a model of bidder participation) that generally have 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.

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

Proof of Proposition 1. Let i be a logger and j a mill. Given an entry profile p , let $P(l, m)$ denote the probability that of the bidders $k \neq i, j$, exactly l loggers and m mills enter. Then

$$\Pi_i^\tau(p) = \sum_{n_L, n_M} \{\pi_L^\tau(l+1, m+1)p_j + \pi_L^\tau(l+1, m)(1-p_j)\} P(l, m) \quad (10)$$

and

$$\Pi_j^\tau(p) = \sum_{l, m} \{\pi_M^\tau(l+1, m+1)p_i + \pi_M^\tau(l, m+1)(1-p_i)\} P(l, m) \quad (11)$$

From Li and Riley (1999), the bracketed term in (10) is no greater than $\pi_L^\tau(l+1, m)$, while the bracketed term in (11) is no less than $\pi_M^\tau(l+1, m+1)$. So $\Pi_i^\tau(p) \leq \sum_{l, m} \pi_L^\tau(l+1, m)P(l, m)$ and $\Pi_j^\tau(p) \leq \sum_{l, m} \pi_M^\tau(l+1, m+1)P(l, m)$. Therefore the assumed condition implies that $\Pi_j^s(p) > \Pi_i^s(p)$ for any logger i and mill j and entry profile p . Moreover, Maskin and Riley's (2000) results imply that for any logger i and mill j and entry profile p , $\Pi_i^s(p) \geq \Pi_i^o(p)$ and $\Pi_j^o(p) \geq \Pi_j^s(p)$, so in addition $\Pi_j^o(p) > \Pi_i^o(p)$. It follows that in any entry equilibrium, if some logger enters with positive probability, then every mill strictly prefers to enter and will enter with probability one. The remaining argument is straightforward. *Q.E.D.*

Proof of Proposition 2. The proof makes use of two key facts arising from the analysis of Maskin and Riley (2000) and Li and Riley (1999). First, for any entry strategies p , $\Pi_i^s(p) \geq \Pi_i^o(p)$ for any logger i and $\Pi_j^s(K) \leq \Pi_j^o(p)$ for any mill j . Second, for either auction format $\tau \in \{o, s\}$ and any bidder i , $\Pi_i^\tau(p)$ is decreasing in p .

For a given vector of type-symmetric entry strategies p , let p_L and p_M denote the entry probabilities of loggers and mills, and $\Pi_L^\tau(p_L, p_M)$, $\Pi_M^\tau(p_L, p_M)$ their expected profits from entry. Fix an auction format τ . From above, if (p_L, p_M) and (p'_L, p'_M) are both type-symmetric entry equilibria, and $p'_M > p_M$, then $p'_L < p_L$. So among type-symmetric entry equilibria, there is one with the most mill entry and least logger entry. Finding this equilibrium is straightforward. If $\Pi_L^\tau(0, 1) < K$, find the unique equilibrium with $p_L = 0$ and $p_M \geq 0$. If $\Pi_L^\tau(0, 1) \geq K$, find the unique equilibrium with $p_L \geq 0$ and $p_M = 1$.

Using the first fact above, it is straightforward to check that the type-symmetric open auction entry equilibrium with the most mill entry and least logger entry will have more mill entry and less logger entry than the type-symmetric sealed auction equilibrium with the most mill entry and least logger entry. This proves the result. *Q.E.D.*

Proof of Proposition 3. Let $\Pi_i^c(p)$ denote the profits of bidder i from entering if mills collude, and similarly for type-symmetric entry profiles define $\Pi_L^c(p_L, p_M)$ and $\Pi_M^c(p_L, p_M)$ as expected bidder profits. We have $\Pi_L^c(p_L, p_M) = \Pi_L^o(p_L, p_M)$ and $\Pi_M^c(p_L, p_M) \geq \Pi_M^o(p_L, p_M)$. Moreover, $\Pi_i^c(p)$ is decreasing in p for any bidder i . Therefore we can use precisely the argument from the above proof to show that the type-symmetric collusive open auction entry equilibrium with the most mill entry and least logger entry will have more mill entry and less logger entry than the corresponding type-symmetric competitive open auction entry equilibrium. *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.²⁴

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 (8), 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: $v_{it} = \exp(X_t \beta_X + N_t \beta_N) \cdot \xi_{it}$. Let $\bar{X} = \mathbb{E}_{X_t}[X_t]$ and $\bar{N} = \mathbb{E}_{N_t}[N_t]$. We assume that for the “stronger” bidder type in a given auction (i.e. mills if any are present, otherwise loggers) $\exp(\bar{X} \beta_X + \bar{N} \beta_N) \cdot \xi_{it} \leq 500$, so that for the average tract in our sample, the highest possible value is \$500 per mbf. This assumption implies an upper bound on the value distribution $\bar{v}_t(X_t, u_t, N_t)$:

$$\bar{v}(X_t, u_t, N_t) = 500 \cdot \frac{\exp(X_t \beta_X + N_t \beta_N)}{\exp(\bar{X}_t \beta_X + \bar{N}_t \beta_N)}.$$

²⁴An 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.

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 2, 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,” *Econometrica*, January 2004, 74(1), pp. 235–267.
- Andrews, D.** “A Conditional Kolmogorov Test,” *Econometrica*, September 1997, 65(5), pp. 1097-1128.
- Asker, J.** “A Study of the Internal Organization of a Bidding Cartel,” NYU Working Paper, 2008.
- 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.
- Athey, S., J. Levin and E. Seira**, “Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions,” Stanford Working Paper 2004.
- Bajari, P.**, “The First-Price Auction with Asymmetric Bidders: Theory and Applications,” Ph.D. Dissertation (University of Minnesota), 1997.
- Bajari, P. and A. Hortacsu**, “The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions,” *Rand Journal of Economics*, Summer 2003, 34(2), pp. 329-355.
- Bajari, P. and L. Ye**, “Deciding Between Competition and Collusion,” *The Review of Economics and Statistics*, November 2003, pp. 971-989.
- Baldwin, L., R. Marshall, and J.-F. Richard.**, “Bidder Collusion at U.S. Forest Service Timber Auctions,” *Journal of Political Economy*, August 1997, 105(4), pp. 657–699.
- Berry, S.**, “Estimation of a Model of Entry in the Airline Industry,” *Econometrica*, July 1992, 60(4), pp. 889–917.
- Branman, L., and L. Froeb**, “Mergers, Cartels, Set-Asides and Bidding Preferences in Asymmetric Second-price Auctions,” *Review of Economics and Statistics*, 2000, 82(2), 283-290
- Brendstrup, B. and H. Paarsch**, “Nonparametric Estimation of Dutch and First-Price, Sealed-Bid Auction Models with Asymmetric Bidders.” Working Paper, University of Iowa Department of Economics, 2003.
- Bresnahan, T. and P. Reiss**, “Entry in Monopoly Markets,” *Review of Economic Studies*, October 1990, 57(4), pp. 531-553.
- Guerre, E., I. Perrigne and Q. Vuong**, “Optimal Nonparametric Estimation of First Price Auctions,” *Econometrica*, May 2000, 68(3), pp. 525–574.

- Froeb, L. and P. McAfee**, “Deterring Bid Rigging in Forest Service Timber Auctions,” U.S. Department of Justice Economic Analysis Group Working Paper, 1988.
- Haile, P.**, “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, June 2001, 92(3), pp. 399–427.
- Haile, P., H. Hong and M. Shum**, “Nonparametric Tests for Common Values in First-Price Sealed Bid Auctions,” Yale University Working Paper, 2003.
- Haile, P. and E. Tamer**, “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy* 2003, 111, pp. 1–52.
- Hansen, R.**, “Sealed-Bid Versus Open Auctions: The Evidence,” *Economic Inquiry*, January 1986, 24(1), pp. 125–142.
- Jofre-Bonet, M. and M. Pesendorfer**, “Estimation of a Dynamic Auction Game,” *Econometrica*, 2003, 71, pp. 1443–1489.
- Johnson, R.**, “Oral Auctions versus Sealed Bids: An Empirical Investigation,” *Natural Resources Journal*, April 1979, 19(2), pp. 315–335.
- Klemperer, P.** *Auctions: Theory and Practice*, Princeton University Press, 2004.
- Krasnokutskaya, E.**, “Identification and Estimation in Highway Procurement Auctions with Unobserved Heterogeneity,” U. Pennsylvania Working Paper, 2004.
- Krasnokutskaya, E. and K. Seim**, “Preferential treatment Program and Participation Decisions in Highway Procurement,” U. Pennsylvania Working Paper, 2005.
- Li, H. and J. Riley**, “Auction Choice,” UCLA Working Paper, 1999.
- Li, T.**, “Econometrics of First-Price Auctions with Entry and Binding Reservation Prices,” *Journal of Econometrics*, May 2005, 126, pp. 173–200.
- Li, T. and X. Zheng**, “Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions,” Vanderbilt Working Paper, 2006.
- Matthews, S.** “Comparing Auctions for Risk Averse Buyers: A Buyer’s Point of View,” *Econometrica*, May 1987, 55(3), pp. 633–646.
- Maskin, E. and J. Riley**, “Asymmetric Auctions,” *Review of Economic Studies*, July 2000, 67(3), pp. 413–438.
- Mead, W.**, *Competition and Oligopsony in the Douglas Fir Timber Industry*, Berkeley, CA: University of California Press, 1966.
- Perrigne, I.**, “Random Reserve Prices and Risk Aversion in Timber Sale Auctions,” Mimeo, Pennsylvania State University, 2003.
- Pesendorfer, M.**, “A Study of Collusion in First-Price Auctions,” *Review of Economic Studies*, 2000, 67, pp. 381–411.

- Porter, R. and J. D. Zona**, “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 101, 1993.
- Porter, R. and J. D. Zona**, “Ohio School Milk Markets: An Analysis of Bidding,” *RAND Journal of Economics*, 30, 1999.
- Schuster, E. and M. Niccolucci**, “Sealed-Bid Versus Oral-Auction Timber Offerings: Implications of Imperfect Data,” *Canadian Journal of Forest Research*, 1994, 24, pp. 87–91.
- Stone, J. and D. Rideout**, “Does Sealed Bidding Promote Competition in the Auctioning of Public Timber?” *Journal of Forest Economics*, 1997, 3(2), pp. 133–142.
- United States Forest Service**, *Forest Service Handbook*, various sections and editions.
- United States House of Representatives**, “Timber Sale Methods and Related Practices,” Hearings before the Subcommittee on Forests, Committee on Agriculture, on H.R. 5863, February 7-9 and April 5, 1977.
- United States Senate**, “Timber Sale Bidding Procedures,” Hearings before the Subcommittee on Public Lands and Resources, Committee on Energy and Natural Resources, on S. 1360, May 18, 1977.
- Vickrey, W.**, “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *Journal of Finance*, 1961, 16(1), pp. 8-37.

Table 1A: Summary Statistics for Northern Sales

N	Open Auctions				Sealed Auctions			
	Full Sample		Selected		Full Sample		Selected	
	886		732		347		339	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
<i>Auction Outcomes</i>								
Winning Bid (\$/mbf)	61.82	43.84	63.27	44.91	69.45	46.64	69.42	46.63
Entrants	4.07	2.41	4.13	2.44	4.37	2.81	4.40	2.83
# Loggers Entering	2.53	2.43	2.75	2.46	3.18	2.57	3.23	2.57
# Mills Entering	1.53	1.70	1.38	1.68	1.19	1.74	1.17	1.74
Fraction Loggers Entering	0.59	0.40	0.64	0.39	0.75	0.34	0.76	0.33
Logger Wins Auction	0.54	0.50	0.59	0.49	0.67	0.47	0.68	0.47
<i>Appraisal Variables</i>								
Volume of timber (hundred mbf)	32.00	43.06	23.02	34.30	16.70	30.01	14.65	25.50
Reserve Price (\$/mbf)	24.70	24.66	25.68	25.46	26.65	24.30	26.64	24.44
Selling Value (\$/mbf)	252.60	131.88	253.04	130.67	259.44	125.05	259.35	125.33
Road Construction (\$/mbf)	5.86	9.57	4.36	8.69	2.94	7.61	2.71	7.44
No Road Construction	0.59	0.49	0.68	0.47	0.78	0.42	0.79	0.41
Logging Costs (\$/mbf)	79.87	64.40	78.61	64.49	79.86	63.42	79.60	63.61
Manufacturing Costs (\$/mbf)	108.82	85.81	107.53	86.09	112.75	87.03	112.65	87.38
<i>Sale Characteristics</i>								
Contract Length (months)	23.02	17.93	22.19	16.35	16.78	14.72	15.94	13.38
Species Herfindal	0.61	0.28	0.61	0.28	0.59	0.27	0.59	0.27
Density of Timber (mbf/acres)	7.83	7.01	7.85	7.20	8.91	8.21	8.97	8.26
Salvage Sale	0.37	0.48	0.37	0.48	0.40	0.49	0.41	0.49
Scale Sale	0.41	0.49	0.38	0.49	0.37	0.48	0.36	0.48
Quarter of Sale	2.43	0.99	2.45	0.99	2.47	0.98	2.47	0.98
Year of Sale	86.32	2.41	86.32	2.45	85.95	2.61	85.94	2.61
Housing Starts	1557.53	255.73	1572.33	235.52	1542.04	274.69	1540.41	275.81
<i>Potential Competition</i>								
Logging companies in county	44.41	20.88	43.17	21.32	41.48	22.00	41.73	22.10
Sawmills in County	8.67	4.32	8.45	4.35	7.67	4.26	7.56	4.15
Active Loggers (active in District in prior 12 months)	20.14	9.71	19.91	9.71	19.24	8.90	19.47	8.86
Active Manufacturers (active in District in prior 12 months)	5.20	2.14	5.25	2.20	5.77	2.47	5.79	2.48

Table 1B: Summary Statistics for California Sales

N	Open Auctions				Sealed Auctions			
	Full Sample 1290		Selected 325		Full Sample 774		Selected 382	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Auction Outcomes								
Winning Bid (\$/mbf)	91.23	143.40	85.10	102.93	79.25	61.90	80.39	63.20
Entrants	4.08	2.35	3.72	2.24	3.74	2.57	4.27	2.76
# Loggers Entering	1.33	1.58	1.95	1.89	2.81	2.22	3.01	2.35
# Mills Entering	2.75	1.90	1.77	1.77	0.94	1.41	1.26	1.57
Fraction Loggers Entering	0.32	0.33	0.55	0.39	0.78	0.31	0.73	0.31
Logger Wins Auction	0.26	0.44	0.50	0.50	0.74	0.44	0.67	0.47
Appraisal Variables								
Volume of timber (hundred mbf)	67.26	50.41	23.77	24.14	9.22	17.68	11.24	12.45
Reserve Price (\$/mbf)	36.79	36.72	37.31	49.64	36.07	32.48	33.26	32.28
Selling Value (\$/mbf)	270.83	99.48	233.41	143.15	246.92	251.63	239.72	119.74
Road Construction (\$/mbf)	9.67	12.58	4.39	9.76	1.09	4.31	1.58	5.23
No Road Construction	0.32	0.46	0.64	0.48	0.90	0.29	0.87	0.34
Logging Costs (\$/mbf)	108.28	44.78	88.42	58.59	86.56	56.74	97.73	58.49
Manufacturing Costs (\$/mbf)	122.74	42.85	100.02	61.11	99.62	62.89	104.93	60.28
Sale Characteristics								
Contract Length (months)	27.54	14.57	16.23	10.34	10.24	7.46	11.97	6.78
Species Herfindal	0.56	0.23	0.59	0.25	0.60	0.24	0.61	0.25
Density of Timber (mbf/acres)	11.54	14.75	11.93	16.94	18.38	220.16	11.37	15.83
Salvage Sale	0.13	0.34	0.24	0.43	0.36	0.48	0.25	0.43
Scale Sale	0.90	0.30	0.75	0.44	0.66	0.47	0.75	0.43
Quarter of Sale	2.37	1.02	2.58	1.00	2.71	0.89	2.62	0.96
Year of Sale	85.26	2.12	85.58	2.29	85.61	2.29	85.07	2.15
Housing Starts	1593.72	254.02	1564.30	235.05	1559.08	247.12	1578.31	264.04
Potential Competition								
Logging companies in county	21.76	18.30	20.75	18.66	20.00	17.39	21.09	19.00
Sawmills in County	6.28	6.05	5.77	5.08	6.00	6.13	6.70	7.45
Active Loggers (active in Forest in prior 12 months)	10.32	7.08	10.11	6.86	10.48	5.98	11.29	6.62
Active Manufacturers (active in Forest in prior 12 months)	5.99	3.34	5.42	3.54	5.26	2.76	5.48	2.85

Table 2: Choice of Sale Method

<i>Dependent Variable: Dummy if auction is sealed bid (Logit regression)</i>				
	(1)		(2)	
	Northern		California	
	coefficient	s.e.	coefficient	s.e.
Appraisal Controls				
Ln(Reserve Price)	-0.063	(0.115)	-0.157	(0.145)
Ln(Selling Value)	0.085	(0.110)	-0.129	(0.111)
Ln(Logging Costs)	0.060	(0.468)	-1.159	(0.491)
Ln(Manufacturing Costs)	0.230	(0.717)	0.343	(0.158)
Ln(Road Costs)	-0.176	(0.086)	-0.129	(0.100)
Other Sale Characteristics				
ln(Contract Length/volume)	-0.756	(4.726)	-6.943	(6.317)
Species Herfindal	-0.211	(0.434)	-0.590	(0.427)
Density of Timber (hmbf/acres)	-1.185	(1.061)	0.042	(0.107)
Salvage Sale (Dummy)	0.112	(0.176)	0.099	(0.242)
Scale Sale (Dummy)	0.328	(0.198)	-0.729	(0.269)
ln(Monthly US House Starts)	-1.168	(1.038)	-4.216	(1.226)
Volume Controls (Dummy Variables):				
Volume: 1.5-3 hundred mbf	-0.115	(0.339)	-1.685	(0.600)
Volume: 3-5	-0.327	(0.360)	-2.262	(0.631)
Volume: 5-8	-0.411	(0.382)	-2.728	(0.657)
Volume: 8-12	-0.744	(0.410)	-3.458	(0.688)
Volume: 12-20	-0.847	(0.405)	-3.832	(0.700)
Volume: 20-40	-1.477	(0.461)	-7.045	(0.762)
Volume: 40-65	-1.758	(0.517)	-8.085	(0.802)
Volume: 65-90	-1.378	(0.550)	-8.796	(0.877)
Volume: 90+	-2.479	(0.583)	-9.833	(0.930)
Potential Competition				
ln(Loggers in County)	-0.089	(0.277)	0.972	(0.214)
ln(Sawmills in County)	0.254	(0.358)	-1.048	(0.279)
ln(Active Loggers)	0.212	(0.177)	0.189	(0.111)
ln(Active Manufacturers)	-0.082	(0.120)	0.164	(0.095)
Additional Controls (Dummy Variables)				
<i>Chi-Squared Statistics (p-value in parenthesis)</i>				
Years	35.96	(0.005)	68.01	(0.000)
Quarters	4.71	(0.195)	4.48	(0.214)
Species	14.67	(0.401)	12.59	(0.127)
Location	114.77	(0.000)	139.96	(0.000)
	N=1233		N=2064	
	LR chi2 (68)	283.24	LR chi2 (55)	1801.07
	P-value	0.000	P-value	0.000
	Pseudo-R2	0.19	Pseudo-R2	0.66

Table 3: Effect of Auction Method on Sale Outcomes

<i>Dependent Variable:</i>	(1) ln(Logger Entry)	(2) ln(Mill Entry)	(3) Loggers/Entrants	(4) Logger Wins	(5) ln(Price)	(6) ln(Price) ¹
Panel A: Northern Sales (N= 1071 Sales)						
<i>Regression with No Interactions Between Sealed and Covariates²</i>						
Sealed Bid Effect	0.089 (0.036)	-0.014 (0.030)	0.056 (0.016)	0.039 (0.026)	0.094 (0.038)	0.055 (0.032)
<i>Regression with Interactions Between Sealed and All Covariates</i>						
Sealed Bid Effect on Sample	0.097 (0.036)	-0.010 (0.031)	0.058 (0.016)	0.038 (0.027)	0.099 (0.039)	0.060 (0.033)
<i>Matching Estimate³</i>						
Sealed Bid Effect on Sample	0.100 (0.048)	0.018 (0.053)	0.052 (0.029)	0.034 (0.039)	0.118 (0.064)	0.091 (0.055)
Panel B: California Sales (N= 707 Sales)						
<i>Regression with No Interactions Between Sealed and Covariates²</i>						
Sealed Bid Effect	0.101 (0.045)	-0.026 (0.038)	0.058 (0.020)	0.036 (0.036)	0.027 (0.051)	-0.026 (0.040)
<i>Regression with Interactions Between Sealed and All Covariates</i>						
Sealed Bid Effect on Sample	0.099 (0.044)	-0.022 (0.038)	0.056 (0.020)	0.035 (0.035)	0.026 (0.050)	-0.037 (0.039)
<i>Matching Estimate³</i>						
Sealed Bid Effect on Sample	0.106 (0.062)	-0.123 (0.067)	0.097 (0.034)	0.107 (0.051)	-0.038 (0.127)	0.005 (0.087)

Notes: Robust standard errors in parentheses, matching standard errors computed following Abadie and Imbens (2006)

1. Specification includes number of entering mills and loggers in addition to sale controls.

2. See Appendix Tables 1A and 2A for full set of controls and coefficients.

3. Number of matches = 4 using the estimated propensity score.

Table 4: Bid and Entry Distributions for Sealed Auctions

	Panel A: Northern Sales				Panel B: California			
	(1)		(2)		(1)		(2)	
	Bid Distribution (Weibull)		Logger Entry (Binomial)		Bid Distribution (Weibull)		Logger Entry (Binomial)	
	coeff.	s.e.	coeff.	s.e.	coeff.	s.e.	coeff.	s.e.
	<i>ln(λ)</i>		<i>α</i>		<i>ln(λ)</i>		<i>α</i>	
Ln(Reserve Price)	0.404	(0.034)	-0.314	(0.057)	0.602	(0.037)	-0.388	(0.061)
Ln(Selling Value)	-0.037	(0.027)	0.025	(0.068)	-0.021	(0.024)	-0.037	(0.051)
Ln(Manufacturing Costs)	1.034	(0.180)	1.255	(0.318)	0.014	(0.032)	0.194	(0.065)
Ln(Logging Costs)	-0.480	(0.164)	-0.934	(0.234)	-0.209	(0.109)	-1.401	(0.206)
Ln(Road Costs)	0.001	(0.025)	-0.148	(0.047)	-0.013	(0.023)	-0.175	(0.062)
Species Herfindal	-0.124	(0.103)	-0.301	(0.176)	-0.200	(0.084)	-0.445	(0.170)
Density of Timber (hmbf/acres)	-0.009	(0.003)	-0.006	(0.005)	-0.003	(0.002)	-0.001	(0.003)
Salvage Sale (Dummy)	-0.005	(0.045)	-0.011	(0.080)	-0.019	(0.049)	-0.417	(0.102)
Scale Sale (Dummy)	-0.008	(0.053)	-0.161	(0.090)	0.133	(0.056)	0.219	(0.104)
Ln(Volume)	-0.063	(0.028)	-0.284	(0.051)	-0.041	(0.032)	-0.307	(0.063)
No Mill Entrants (Dummy)	-0.120	(0.068)	-0.460	(0.117)	-0.089	(0.059)	-0.433	(0.120)
Min(Mill Entrants,5)	0.091	(0.022)	-0.059	(0.045)	0.072	(0.020)	-0.053	(0.047)
Active Loggers			-0.041	(0.006)			-0.050	(0.007)
Min(Logger Entrants,5)	0.034	(0.016)			0.074	(0.014)		
Mill (Dummy)	0.284	(0.032)			0.160	(0.026)		
Mill (Dummy) * (Mill Entrants=1)	-0.109	(0.073)			-0.128	(0.052)		
Additional Controls	Forest-District, Year, Species Dummies		Forest-District, Year, Species Dummies		Forest, Year, Species Dummies		Forest, Year, Species Dummies	
	<i>ln(p)</i>				<i>ln(p)</i>			
No Mill Entrants (Dummy)	-0.081	(0.070)			-0.202	(0.070)		
Min(Mill Entrants,5)	0.015	(0.022)			-0.009	(0.022)		
Min(Logger Entrants,5)	0.016	(0.018)			-0.005	(0.015)		
Mill(Dummy)	0.074	(0.066)			-0.031	(0.052)		
Mill (Dummy) * (Mill Entrants=1)	-0.386	(0.122)			-0.076	(0.100)		
Constant	1.134	(0.103)			1.329	(0.090)		
	<i>ln(θ)</i>				<i>ln(θ)</i>			
Constant	-0.502	(0.120)			-0.416	(0.124)		
	N=1421		N = 339		N=1565		N = 382	

Note: Bid distribution estimated from sales with two or more bidders.

Table 5: Actual Outcomes vs. Outcomes Predicted by Model

		(1)	(2)	(3)		
	N	Actual	Predicted (bidding only)	Predicted (bidding + entry)		
Panel A: Northern Sales						
Sealed Bid Sales						
Avg. Bid	1492	59.6	58.2	(1.4)	57.4	(1.3)
Avg. Logger Bid	1096	50.8	48.7	(1.4)	47.4	(1.4)
Avg. Mill Bid	396	83.8	84.7	(2.7)	85.2	(2.7)
Avg. Sale Price (\$/mbf)	339	69.4	69.9	(1.4)	70.4	(1.6)
Avg. Revenue (\$000s)	339	111.4	108.1	(4)	109.9	(4.2)
% Sales won by Loggers	339	68.1	68.0	(0.90)	65.0	(0.01)
Avg. Logger Entry	339	3.23			3.23	(0.09)
Open Auction Sales						
Avg. Sale Price (Competition)	732	63.3	67.9	(1.8)	67.8	(2.1)
Avg. Sale Price (Collusion)	732	63.3	44.2	(1.3)	44.1	(2.2)
Avg. Revenue (Competition)	732	144.7	152.7	(6.8)	154.8	(7.9)
Avg. Revenue (Collusion)	732	144.7	61.0	(2)	64.7	(5.0)
% Sales won by Loggers	732	59.0	56.0	(0.01)	54.4	(0.02)
Avg. Logger Entry	732	2.75			2.67	(0.17)
Panel B: California Sales						
Sealed Bid Sales						
Avg. Bid	1630	73.6	74.7	(2.3)	74.2	(2.3)
Avg. Logger Bid	1150	64.0	63.6	(2.1)	62.3	(2.4)
Avg. Mill Bid	480	96.5	101.2	(3.5)	102.8	(3.8)
Avg. Sale Price (\$/mbf)	382	80.4	83.8	(2.1)	84.4	(2.4)
Avg. Revenue (\$000s)	382	103.1	110.7	(3.8)	111.9	(4.0)
% Sales won by Loggers	382	66.8	66.4	(1.2)	62.6	(1.3)
Avg. Logger Entry	382	3.01			3.01	(0.07)
Open Auction Sales						
Avg. Sale Price (Competition)	325	85.1	87.2	(2.7)	86.7	(3.1)
Avg. Sale Price (Collusion)	325	85.1	46.1	(1.2)	51.0	(1.6)
Avg. Revenue (Competition)	325	227.0	244.7	(9.7)	242.4	(10.9)
Avg. Revenue (Collusion)	325	227.0	93.2	(2.6)	112.9	(5.6)
% Sales won by Loggers	325	50.5	48.2	(1.1)	43.6	(1.8)
Avg. Logger Entry	325	1.95			1.90	(0.13)

Table 6: Actual versus Predicted Sale Prices by Mill Participation

		(1)	(2)	(3)
	N	Actual	Predicted (bidding only)	Predicted (bidding + entry)
Zero Mills				
Sealed Bid Sales	181	51.7	51.4	51.4
Open Auction Sales	321	49.8	50.5	47.1
One Mill				
Sealed Bid Sales	70	66.8	64.6	66.9
Open Auction Sales	150	50.0	52.2	59.5
Two or More Mills				
Sealed Bid Sales	88	108.1	112.1	112.2
Open Auction Sales	261	87.5	98.5	98.0

Note: Average sale prices are for Northern region tracts.

Table 7: Welfare Effects of Sealed vs. Open Auctions

	(1) Sealed	(2) Open (Comp.)	(3) Sealed - Open Difference	(4) Open (Part. Coll.)	(5) Sealed - Open Difference
Panel A: Northern Sales (N=1071)					
Exogenous Entry					
Avg. Sale Price (\$/mbf)	68.56	68.53	0.03	(0.04)	64.58 3.98 (0.24)
Avg. Sale Revenue (\$000s)	138.33	138.64	-0.32	(0.06)	124.33 14.00 (1.22)
Avg. Sale Surplus (\$000s)	176.99	177.06	-0.08	(0.02)	
% Sales Won by Loggers	60.14	59.68	0.46	(0.162)	
Endogenous Entry					
Avg. Sale Price (\$/mbf)	69.68	68.19	1.49	(0.71)	64.11 5.57 (0.82)
Avg. Sale Revenue (\$000s)	143.04	139.87	3.17	(2.24)	125.32 17.72 (2.74)
Avg. Sale Surplus (\$000s)	156.53	156.84	-0.30	(2.70)	
% Sales Won by Loggers	59.65	57.19	2.46	(0.00)	
Logger Entry	3.10	2.76	0.34	(0.12)	
Panel B: California Sales (N = 707)					
Exogenous Entry					
Avg. Sale Price (\$/mbf)	85.53	85.20	0.32	(0.07)	79.39 6.13 (0.39)
Avg. Sale Revenue (\$000s)	172.64	172.10	0.55	(0.13)	154.43 18.21 (1.36)
Avg. Sale Surplus (\$000s)	203.61	203.65	-0.04	(0.02)	
% Sales Won by Loggers	58.21	57.81	0.40	(0.002)	
Endogenous Entry					
Avg. Sale Price (\$/mbf)	88.36	83.76	4.60	(2.26)	78.55 9.81 (2.21)
Avg. Sale Revenue (\$000s)	182.38	168.20	14.18	(5.43)	152.83 29.55 (5.39)
Avg. Sale Surplus (\$000s)	181.82	170.58	11.24	(6.86)	
% Sales Won by Loggers	55.84	52.66	3.18	(0.01)	
Logger Entry	2.72	2.35	0.37	(0.14)	

Note: Each entry is an average prediction over all tracts in the sample. Bootstrap standard errors in parentheses.

Figure 1A: Actual vs. Estimated Density of Sealed Bid Residuals (Northern Sales)




Figure 1B: Actual vs. Estimated Density of Sealed Bid Residuals (CA Sales)




Figure 2A: Estimated Value Distributions and Bid Functions for the Case of Two Loggers and Two Mills (Northern Sales)




Figure 2B: Estimated Value Distributions and Bid Functions for the Case of Two Loggers and Two Mills (CA Sales)

