

When Should Sellers Use Auctions?

James W. Roberts* Andrew Sweeting†

June 29, 2011

Abstract

A bidding process can be organized so that offers are submitted simultaneously or sequentially. In the latter case, potential buyers can condition their behavior on previous entrants' decisions. The relative performance of these mechanisms is investigated when entry is costly and selective, meaning that potential buyers with higher values are more likely to participate. A simple sequential mechanism can give both buyers and sellers significantly higher payoffs than the commonly used simultaneous bid auction. The findings are illustrated with parameters estimated from simultaneous entry USFS timber auctions where our estimates predict that the sequential mechanism would increase revenue and efficiency.

JEL CODES: D44, L20

Keywords: Auctions, Entry, Selection, Mechanism Design, Mergers and Acquisitions

*Department of Economics, Duke University. Contact: j.roberts@duke.edu.

†Department of Economics, Duke University and NBER. Contact: atsweet@duke.edu.

We are extremely grateful to Susan Athey, Jonathan Levin and Enrique Seira for sharing the data with us and to Coleman Doggett, Kent Lou and Doug MacDonald for their insight on the timber industry. We would also like to thank John Asker, Brendan Daley, Ken Hendricks, Jakub Kastl, Greg Lewis, David McAdams, Brian McManus, Rob Porter, Steve Tadelis and numerous seminar participants for useful feedback. Excellent research assistance was provided by Vivek Bhattacharya. Both authors gratefully acknowledge financial support from Duke University. Any errors are our own.

1 Introduction

The simultaneous bid auction is a standard method for sellers to solicit offers from buyers. A simple alternative is for a seller to ask buyers to make offers sequentially. If it is costly for buyers to participate, the sequential mechanism will tend to be more efficient than the simultaneous auction because later potential buyers can condition their participation decisions on earlier bids. However, the sequential mechanism's greater efficiency may not produce higher revenues because while the possibility of deterring later potential entrants can lead early bidders to bid aggressively, the fact that later firms might be deterred will tend to reduce revenues. The relative revenue performance of the mechanisms will therefore depend on whether the threat of potential future competition, which can raise bids in the sequential mechanism, is more valuable to the seller than actual competition, which will tend to be greater in the simultaneous auction.

The relative revenue performance of these alternative mechanisms has direct implications for how assets should be sold. In the case of how to structure the sale of corporations, this question has attracted attention from practitioners and other commentators since the Delaware Supreme Court's 1986 *Revlon* decision charged a board overseeing the sale of a company with the duty of "getting the best price for the stockholders" (*Revlon v McAndrews & Forbes Holdings* (1986)). In practice, corporate sales occur through a mixture of simultaneous and sequential mechanisms, with sequential mechanisms sometimes taking the form of "go-shop" arrangements where a seller may reach an agreement with one firm while retaining the right to solicit other offers, to which the first firm may be able to respond.^{1,2}

Surprisingly, the only attempt to date to directly address this relative performance question is Bulow and Klemperer (2009) (BK hereafter). They compare the revenue and efficiency performances of the commonly-used simultaneous bid second-price auction with a similarly simple, sequential mechanism. In this second mechanism, buyers are approached in turn, and upon observing the history of offers, each chooses whether to enter and attempt to outbid the current high bidder. If the incumbent is outbid, the new entrant can make a jump bid that may potentially deter later firms from participating. The incumbent at the end of the game pays the standing price. As BK note (see also Subramanian (2008), Wasserstein (2000))³,

¹A "go-shop" clause allows a seller to come to an agreement on an initial price with a buyer and retain the right to solicit bids from other buyers for the next 30-60 days. If a new, higher offer is received, then according to the "match right", which is often included in the agreement with the initial buyer, the seller must negotiate with the first buyer (for 3-5 days, for example) to see if it can match the terms of the new, higher offer.

²There are numerous theory papers, some related directly to the field of corporate finance, that consider sequential mechanisms similar to the one considered here. Examples include Fishman (1988), Daniel and Hirshleifer (1998) and Horner and Sahuguet (2007).

³Boone and Mulherin (2007) describe various sale methods for corporate takeovers and establish that the

these simple mechanisms can be thought of as spanning the range of sale processes that are actually used. In the comparison, BK assume that potential bidders only know the distribution from which values are drawn prior to entering, and have no additional information about their own value. After entry they find out their values for sure. These assumptions are common in the auction literature as they provide greater analytic tractability. Under this informational assumption, together with the assumption that bidders are symmetric and the seller cannot set a reserve price, BK show that “sellers will generally prefer auctions and buyers will generally prefer sequential mechanisms” (p. 1547).

This result holds in BK’s model because, in the equilibrium they consider, early bidders with high enough values submit bids that deter all future potential entry (all future potential entrants have the same beliefs about their values prior to entry), and there is too much deterrence from the seller’s perspective. Thus, he would prefer the greater actual competition provided by the auction. In particular, deterrence means that later potential entrants with high values will not enter, which decreases both the expected value of the winning bidder and the value that an incumbent has to pay. In contrast, buyers prefer the sequential mechanism as expenditures on entry costs are lower. This effect also tends to increase social efficiency.

In light of their result, BK interpret the use of sequential mechanisms as evidence that buyer’s preferences can determine the choice of mechanism, consistent with the fact that some influential buyers, such as Warren Buffett, have explicit policies that they will not “waste time” by participating in auctions.

In this paper, we consider a similar comparison, except that we extend BK’s model to allow potential buyers to receive a noisy signal about their valuation prior to deciding whether to enter either mechanism. After entry, they find out their values for sure, as in BK’s model. This structure results in a “selective entry” model, where firms enter if they receive high enough signals, and firms with higher values are more likely to enter.^{4,5} We believe that this is a natural model to describe settings where firms are likely to have some imperfect information about their value for an asset based on publicly available information, but must conduct costly research to discover additional information that will affect their value.⁶ We

set of potential bidders in these contests is larger than the set submitting public bids.

⁴The precision of the signal determines how selective the entry process is. In its limits, the model can approach the polar cases of (a) perfect selection, which we term the S model after Samuelson (1985), whereby a firm knows its value exactly when taking its entry decision, and (b) no selection, which we term the LS model after Levin and Smith (1994), whereby a firm knows nothing of its value when taking its entry decision.

⁵Selective entry contrasts with standard assumptions in the empirical entry literature (e.g., Berry (1992)) where entrants may differ from non-entering potential entrants in their fixed costs or entry costs, but not along dimensions such as marginal costs or product quality that affect competitiveness or the profits of other firms once they enter.

⁶Examples include oil and gas leases, timber sales, government procurement contracts and firm takeovers. Recently there has been some work allowing for selection in empirical auction research. The dominant way this is done is by assuming that bidders know their value precisely prior to entry, i.e. by assuming perfect

also allow for potential buyers to be asymmetric, which is another important feature of many real-world settings.

Using numerical analysis, which becomes necessary once either asymmetries or selective entry are added to the model, we show that the sequential mechanism can give the seller *higher* expected revenues than the simultaneous auction even when buyers' signals about their values are quite noisy. When the entry process is quite selective and/or entry costs are large, the difference in revenues can be substantial, and, as a comparison, the increase in revenues from using the sequential mechanism is much larger than the returns to using an optimal reserve price in the simultaneous auction. As in BK's analysis, the sequential mechanism is more efficient, and the sequential mechanism generally gives higher expected payoffs to both buyers and sellers. This result would obviously lead to a different interpretation of why sequential mechanisms are sometimes used, and because the sequential mechanism increases the payoffs of buyers, it is still consistent with comments like those of Warren Buffett. Our findings are also consistent with observed differences in target shareholder returns in private equity transactions documented by Subramanian (2008). He compares returns when companies are sold using a go-shop process and a process where many firms are simultaneously asked to submit bids before a winner is selected. He finds that target shareholder returns are 5% higher for go-shops and he argues that, even though go-shop agreements introduce asymmetries between bidders into the sale process, they are preferable for both buyers and sellers.

We illustrate our findings using parameters estimated from a sample of (simultaneous) open outcry US Forest Service (USFS) timber auctions. This setting provides a close match to the information structure assumed in our model as a potential bidder can form a rough-estimate of its value based on tract information published by the USFS and knowledge of its own sales contracts and capabilities, and it is also standard for interested bidders to conduct their own tract surveys ("cruises") prior to bidding. It is also a setting where various auction design tools, such as reserve price policies, have been studied by both academics and practitioners in order to try to raise revenues which have often been regarded as too low.⁷

selection. For example, Li and Zheng (2009) compare estimates from both the LS and S models using data on highway lawn mowing contracts from Texas to understand how potential competition may affect procurement costs, and Li and Zheng (2010) test the LS and S models using timber auctions in Michigan. Marmer, Shneyerov, and Xu (2010) extend this literature by testing whether the Li and Zheng (2009) data is best explained by the LS, S or a more general affiliated signal model. They find support for the S and signal models, and they also estimate a very simple version of their signal model. Finally, Gentry and Li (2010) show how partial identification techniques can be used to construct bounds on the primitives of a signal model.

⁷Some examples of studies of timber auction reserve prices include Mead, Schniepp, and Watson (1981), Paarsch (1997), Haile and Tamer (2003), Li and Perrigne (2003) and Aradillas-Lopez, Gandhi, and Quint (2010). All of these papers assume that entry is not endogenous. Academics have also provided expert advice to government agencies about how to set reserve prices (stumpage rates) for timber (e.g. Athey, Cramton, and Ingraham (2003)). In 2006, Governor Tim Pawlenty of Minnesota commissioned a task force

Timber auctions are also characterized by important asymmetries between potential buyers, with sawmills tending to have systematically higher values than loggers.

Our estimates imply that the entry process into timber auctions is moderately selective, while average entry costs are 2.3% of the average winning bid, which is large enough to prevent many loggers from entering auctions. For the (mean) representative auction in our data, our results imply that using a sequential mechanism (with no reserve) would generate a nine times larger increase in revenues than setting the optimal reserve price (the focus of the existing literature) in the simultaneous auction. We also find that the efficiency gains from using the sequential mechanism are large enough that both the USFS revenues and firm profits can increase. Additionally, loggers (the weaker type) win more often. These results suggest that the sequential mechanism may present the USFS and other procurement agencies with an effective, new alternative to commonly used set-aside programs and bid subsidies for ensuring that a certain fraction of projects are won by a targeted set of bidders.⁸

Why does the sequential mechanism tend to produce higher revenues when entry is selective? The key reason is that selective entry changes the nature of the equilibrium in the sequential mechanism in a way that tends to increase both its relative efficiency and the revenues that the seller can extract. With no selection, BK show that the “pre-emptive bidding [which occurs in equilibrium] is crucial: jump-bidding allows buyers to choose partial-pooling deterrence equilibria which over-deter entry relative to the social optimum” (p. 1546). Introducing any degree of selection into the entry process causes the bidding equilibrium (in the unique equilibrium under the D1 refinement which we focus on) to change so that there is full separation, with bids perfectly revealing the value of the incumbent.⁹ At the same time, a potential entrant will enter if it receives a high enough signal about its value. These changes increase the efficiency of the outcome in the sequential mechanism as higher value incumbents deter more entry with higher value potential entrants being more likely to enter. Unlike in BK’s model, the expected value of the winner can be higher in the sequential mechanism, which increases the rents available to all parties. In addition, the change to a separating equilibrium affects the equilibrium level of jump bids. For some values, this will increase the expected amount that bidders pay, benefiting the seller.

Some comments about the nature of our results are appropriate. First, we do not seek to investigate the performance of the state’s timber sale policies, and its report indicates an openness to considering alternative sales mechanisms as well as different reserve prices (Kilgore, Brown, Coggins, and Pfender (2010)).

⁸The USFS has historically used set-asides and recent work (Athey, Coey, and Levin (2011)) suggests this may come at a substantial revenue and efficiency loss relative to using bid subsidies.

⁹This is correct for values less than the upper limit of the value distribution minus the cost of entry. An upper limit on the value distribution is required for technical reasons but we assume that it is sufficiently high that, for practical purposes, all incumbent values are revealed.

to compare revenues with those from the optimal mechanism. Instead, in the same spirit as BK, we are interested in the relative performance of stylized versions of commonly used sales mechanisms, whereas the seller optimal mechanism, which is not known for a model with imperfectly selective entry (Milgrom (2004)), is likely to involve features, such as side payments or entry fees that are rarely observed in practice, and which might require the seller to have implausibly detailed information.¹⁰ The seller would also need to know this information if he wants to set the optimal reserve price in an auction, and, indeed, an attraction of the simple sequential mechanism that we consider is that the only required information concerns the set of potential entrants who should be approached.¹¹ We show below that if the seller has enough information to set an optimal reserve in the sequential mechanism, he can do even better.

On the other hand, what is known about the optimal mechanism in models with costly entry and either no selection or perfect selection suggests that the optimal mechanism should be sequential, which helps to rationalize our results. For example, Cremer, Spiegel, and Zheng (2009) consider the case with *no* selection and McAfee and McMillan (1988) consider a model where buyers know their values but it is costly for the seller to engage additional buyers. In both cases, the optimal mechanism involves some type of sequential search procedure, which stops when a buyer with a high enough value is identified.¹²

Second, while we characterize the unique equilibrium of each mechanism under standard refinements, our revenue comparisons are numerical in nature. This is a necessary cost of allowing for either a more general model of entry, or bidder asymmetries. Our results show that these features matter because the relative performance of the mechanisms can change even when selection is quite imperfect. The computational approach also allows us to provide a substantive empirical application of our model as selective entry and bidder asymmetries are clear features of our data.

Third, the sequential mechanism can be characterized as a multi-round extension of a standard two-player signaling game where an incumbent bidder can use a jump bid to signal its value to later potential entrants. We contribute to the literature on extensions of two-

¹⁰Our approach is therefore similar to analyses of practical mechanisms in other settings, such as Chu, Leslie, and Sorensen (2011) (bundling), Rogerson (2003) (contracts), McAfee (2002) (nonlinear pricing) and Neeman (2003) (auctions).

¹¹In this sense the sequential mechanism satisfies what has come to be known as the “Wilson doctrine” (Wilson (1987)), which suggests that, from a practical standpoint, we ought to be concerned with mechanisms that do not rely on the seller possessing unrealistically detailed information about buyers.

¹²In this environment, Ye (2007) considers two-stage bidding structures where the seller must choose how many firms pay the entry cost ahead of the first stage. His paper clearly shows how to determine the optimal number of entrants. However, he does not consider a wider range of mechanisms that might allow, for example, the seller to set a reserve price or to decide how many firms should enter only after the first stage bids are submitted.

player signaling games by characterizing the unique sequential equilibrium under standard refinements and providing a straightforward recursive algorithm for calculating equilibrium strategies. We also note that our findings relate to the classic limit pricing result of Milgrom and Roberts (1982). As they show in a two-period, two-firm setting, an incumbent’s incentive to deter a competitor’s entry can benefit consumers through lower prices. We find a similar result that the incentive to deter later potential competitors can benefit a seller through higher prices.

Fourth, we note two differences, beside the introduction of selective entry and bidder asymmetries, between our model and the model considered by BK. First, we assume that the number N of potential entrants is fixed and common knowledge to all players, whereas BK’s model allows for some probability ($0 \leq \rho_j \leq 1$) of a j^{th} potential entrant if there are $j - 1$ potential entrants. As these probabilities may equal 1 for $j < N$, and 0 for $j \geq N$ for any N , our model is a special case of theirs. Our choice reflects the standard practice in the empirical literature, which we want to follow when estimating our model.¹³ Second, when modeling the auction mechanism, we focus on the model where potential buyers make simultaneous entry decisions as well as simultaneous bid choices, whereas BK’s primary focus is on a model where firms make sequential entry decisions before bidding simultaneously. However, in BK’s model “no important result is affected if potential bidders make simultaneous, instead of sequential, entry decisions into the auction” (p.1560). We also give some consideration to a sequential entry, simultaneous bid model, and show that our qualitative results are unchanged. Our choice to focus on simultaneous entry into the auction reflects a desire to reduce the computational burden and, more importantly, the fact that simultaneous entry is the appropriate way to model entry into the auctions in our empirical sample (Athey, Levin, and Seira (forthcoming) also apply a simultaneous entry model (with no selection) to USFS timber auctions).¹⁴

The paper proceeds as follows. Section 2 introduces the models of each mechanism and characterizes the equilibria that we examine. Section 3 compares expected revenue and efficiency from the two mechanisms for wide ranges of parameters, and provides intuition for when the sequential mechanism outperforms the auction. Section 4 describes the empirical setting of USFS timber auctions and explains how we estimate our model. Section 5 presents

¹³Examples of this assumption in the auction literature include Athey, Levin, and Seira (forthcoming) and Li and Zheng (2009). Examples elsewhere in empirical work on entry games include Berry (1992), Seim (2006) and Ciliberto and Tamer (2009).

¹⁴The computational burden in the sequential entry, simultaneous bid auction model arises from the fact that later potential entrants’ equilibrium entry thresholds are a function of the complete history of the game and thresholds in earlier rounds, so that it is necessary to solve for all of the thresholds simultaneously. In contrast, in the sequential mechanism, a potential entrant’s equilibrium threshold only depends on the value of the incumbent which (with any degree of selection) is completely revealed by its jump bid.

the parameter estimates and counterfactual results showing that the USFS could improve its revenues by implementing a sequential mechanism. Section 6 concludes.

2 Model

We now describe the model of firms' values and signals, before describing the mechanisms that we are going to compare.

2.1 A General Entry Model with Selection

Suppose that a seller has one unit of a good to sell and gets a payoff of zero if the good is unsold. There is a set of potential buyers who may be one of $\tau = 1, \dots, \bar{\tau}$ types, with N_τ of type τ . In practice $\bar{\tau} = 2$. Buyers have independent private values (IPV), which can lie on $[0, \bar{V}]$, distributed according to $F_\tau^V(V)$. F_τ^V is continuous and differentiable for all types. In this paper we will assume that the density of V is proportional to the log-normal distribution on $[0, \bar{V}]$, and that \bar{V} is high, so that the density of values at \bar{V} is very small.¹⁵ Before participating in any mechanism, a potential buyer must pay an entry cost K_τ . This entry cost can be interpreted as a combination of research costs necessary to learn one's value and participation/bidding costs. Once it pays K_τ , a potential buyer learns its value. We assume that a firm cannot participate without paying K_τ . However, prior to deciding whether to enter, a bidder receives a private information signal about its value. We focus on the case where the signal of potential buyer i of type τ is given by $s_{i\tau} = v_{i\tau} a_{i\tau}$, where $A_\tau = e^{\varepsilon_\tau}$, $\varepsilon_\tau \sim N(0, \sigma_{\varepsilon_\tau}^2)$ and draws of ε are assumed to be i.i.d. across bidders.

Let $F_\tau^S(s)$ be the unconditional distribution of a bidder's signal and $F_\tau^S(s|v)$ be the distribution conditional on a particular value v . In this model, $\sigma_{\varepsilon_\tau}^2$ controls how much potential buyers know about their values before deciding whether to enter. As $\sigma_{\varepsilon_\tau}^2 \rightarrow \infty$, the model will tend towards the informational assumptions of the Levin and Smith (1994) (LS) model in which pre-entry signals contain no information about values. As $\sigma_{\varepsilon_\tau}^2 \rightarrow 0$, it tends towards the informational assumptions of the Samuelson (1985) (S) model where firms know their values prior to paying an entry cost (which is therefore interpreted as a bid preparation or attendance cost). Intermediate values of $\sigma_{\varepsilon_\tau}^2$, implying that buyers have some idea of their values but have to conduct costly research to learn them for sure, seem plausible for most empirical settings. Having received his signal, a potential buyer forms posterior beliefs about his valuation using Bayes Rule.

¹⁵To be precise, $f^V(v|\theta) = \frac{h(v|\theta)}{\int_0^{\bar{V}} h(x|\theta) dx}$, where $h(v|\theta)$ is the pdf of the log-normal distribution.

2.2 Mechanism 1: Simultaneous Entry Second Price Auction

The first mechanism we consider is a simultaneous entry second price or open outcry auction that we model as a two-stage game. In the first stage all potential buyers simultaneously decide whether to enter the auction (pay K_τ) based on their signal, the number of potential entrants of each type and the auction reserve price. In the second stage, entrants then learn their values and submit bids. We assume that an open outcry auction would give the same outcome as an English button auction, so that the good would be awarded to the firm with the highest value at a price equal to the value of the second highest valued entrant or the reserve price if one is used.¹⁶

Following the literature (e.g. Athey, Levin, and Seira (forthcoming)), we assume that players use strategies that form type-symmetric Bayesian Nash equilibria, where “type-symmetric” means that every player of the same type will use the same strategy. In the auction’s second stage, entrants know their values so it is a dominant strategy for each entrant to bid its value. In the first stage, players take entry decisions based on what they believe about their value given their signal. The (posterior) conditional density $g_\tau(v|s_i)$ that a player of type τ ’s value is v when its signal is s_i is defined via Bayes Rule.

The weights that a player places on its prior and its signal when updating its beliefs about its true value depend on the relative variances of the distribution of values and ε (signal noise), and this will also control the degree of selection. A natural measure of the relative variances is $\frac{\sigma_\varepsilon^2}{\sigma_V^2 + \sigma_\varepsilon^2}$, which we will denote α . If the value distribution were not truncated above, player i ’s (posterior) conditional value distribution would be lognormal with location parameter $\alpha\mu_\tau + (1 - \alpha)\ln(s_i)$ and squared scale parameter $\alpha\sigma_{V_\tau}^2$.

The optimal entry strategy in a type-symmetric equilibrium is a pure-strategy threshold rule where the firm enters if and only if its signal is above a cutoff, S_τ^* . S_τ^* is implicitly defined by the zero-profit condition that the expected profit from entering the auction of a firm with the threshold signal will be equal to the entry cost:

$$\int_R^{\bar{V}} \left[\int_R^v (v - x) h_\tau(x|S_\tau^*, S_{-\tau}^*) dx \right] g_\tau(v|S_\tau^*) dv - K_\tau = 0 \quad (1)$$

where $g_\tau(v|s)$ is defined above, and $h_\tau(x|S_\tau^*, S_{-\tau}^*)$ is the pdf of the highest value of other entering firms (or the reserve price R if no value is higher than the reserve) in the auction, given equilibrium strategies. A pure strategy type-symmetric Bayesian Nash equilibrium exists because optimal entry thresholds for each type are continuous and decreasing in the threshold of the other type.

¹⁶Our estimation procedure does not require that other losing bidders bid up to their values.

With multiple types, there can be multiple equilibria in the entry game when types are similar (for example, in the means of their values) even when we assume that only type-symmetric equilibria are played. As explained in Roberts and Sweeting (2011), we choose to focus on an equilibrium where the type with higher mean values has a lower entry threshold (lower thresholds make entry more likely). This type of equilibrium is intuitively appealing and when firms' reaction functions are S-shaped (reflecting, for example, normal or log-normal value and signal noise distributions) and types only differ in the location parameters of their value distributions (i.e., the scale parameter, signal noise variance and entry costs are the same) then there is exactly one equilibrium of this form.¹⁷ Therefore, we assume that types only differ in the location parameters of their value distributions from now on. Given our focus on this type of equilibrium, solving the model is straightforward: we find the S'^* values that satisfy the zero profit conditions for each type and which satisfy the constraint that $S_1'^* < S_2'^*$, where a type 1 firm is the high type (larger location parameter). It is important to note that the issue of type-symmetric multiple equilibria affects only the auction, not the sequential mechanism.

2.3 Mechanism 2: Sequential Mechanism

As BK and others note, the standard alternative to buyers submitting bids simultaneously is a sequential bid process. Here we describe a very simple sequential bid process like that in BK. Potential buyers are placed in some order (which does not depend on their signals, but may depend on types), and the seller approaches each potential buyer in turn. We will call what happens between the seller's approach to one potential buyer and its approach to the next potential buyer a "round". In the first round, the first potential buyer observes his signal and then decides whether to enter the mechanism and learn his value by paying K . If he enters he can choose to place a 'jump bid' above the reserve price, which we assume to be zero. Given entry, submitting a bid is costless.

In the second round the potential buyer observes his signal, the entry decision of the first buyer and his jump bid, and then decides whether to enter himself. If the first firm did not enter and the second firm does, then the second firm can place a bid in exactly the same way as the first firm would have been able to do had he entered. If both enter, the firms bid against each other in a knockout button auction until one firm drops out, in which case it can never return to the mechanism. The remaining firm then has an opportunity to submit

¹⁷We have also estimated the model using a nested pseudo-likelihood procedure which does not require us to use an equilibrium selection rule. The parameter estimates in this case indicate that the difference in mean values between our two types (sawmills and logging companies) are so large that multiple equilibria cannot be supported.

an additional higher jump bid above the bid at which the other firm dropped out. If the second firm does not enter, but the first firm did, then the first firm can either keep its initial bid or submit a higher jump bid.

This procedure is then repeated for each remaining potential buyer, so that in each round there is at most one incumbent bidder and one potential entrant. The complete history of the game (entry decisions and bids, but not signals) is observed by all players. If a firm drops out, or chooses not enter, it is assumed to be unable to re-enter at a later date. The good is allocated to the last remaining bidder at a price equal to the current bid.

A strategy in the sequential model consists of an entry rule and a bidding rule as a function of the round, the potential buyer's signal and value (for bidding) and the observed history. When a potential buyer is bidding against an active opponent in the knockout auction, the dominant strategy is to bid up to its value, so that the firm with the lower value will drop out at a price equal to its value. This does not depend on the selective entry model because values are known at this stage. However, the strategies that firms use to determine their jump bids and entry decisions do depend on selective entry. To place our equilibrium in context, we begin describing what happens when there are no signals and symmetric firms, which are the assumptions made by BK.

Before describing this mechanism's equilibrium, we note that it is straightforward for a seller to implement this mechanism. In particular, the seller needs only to identify potential entrants, specify and commit to a buyer order and establish a program for collecting and distributing information on the entry and bidding behavior of all firms. For sellers that will implement the mechanism many times, such as the USFS, any fixed costs involved in setting up the mechanism should be relatively small.

2.3.1 Equilibrium with No Pre-Entry Signals

Assuming symmetric firms and no pre-entry signals, BK show that any entering firm that learns its value is below some endogenously determined V^S will keep the existing standing bid, while firms with values above V^S will submit a jump bid that deters all future entry, no matter how many rounds are left. This is because all future potential entrants have identical information about their values prior to taking entry decisions. V^S is independent of the round of the game and history to that point and it is determined by the condition that future potential entrants should be indifferent to entering when the incumbent firm's value is above V^S . The deterring bid is determined by the condition that the bidder with a value V^S is indifferent between deterring future entry with this deterring bid and accommodating entry by keeping the standing bid. Thus while in any round all firms with values above V^S submit the same deterring bid, this bid may depend on the round and history of the game.

Equilibrium with no signals is therefore characterized by entry in every round until a firm with a value greater than V^S participates, in which case entry ceases forever. BK show that while this leads to higher expected efficiency than the auction, from the seller's standpoint, too many bidders are deterred from participating and in equilibrium revenues tend to be lower than in the auction.

2.3.2 Equilibrium with Pre-Entry Signals

There are important changes to the nature of the equilibrium when potential buyers receive pre-entry signals. We begin by describing the equilibrium we consider, before explaining the refinements that lead us to focus on it.

A potential entrant in any round n participates if and only if his signal exceeds some threshold $S'_n(v)$, at which the expected profits from entering are zero and is a function of the round, his beliefs about the current incumbent's (if there is one) value (v) and the expected behavior of future potential entrants. Upon entry, an incumbent and a new entrant bid up to their values in the knockout auction. The winner of the knockout auction may then submit a jump bid that may deter future entry. For a bidder with values on $[0, \bar{V} - K]$, its jump bid will perfectly reveal its value and so we assume that a new incumbent jump bids the first time he is able and after placing one jump bid he will not do so again. Therefore, given a bidding function in round n , $\beta(v, \hat{b}_n, n)$, which depends on the bidder's value (v), the standing bid prior to the jump bid being placed (\hat{b}_n - this will be zero when the bidder is the first entrant and otherwise it will be the previous incumbent's value since they will have just lost in a knockout auction prior to a new jump bid being placed) and the round, an incumbent with value v must decide which v' 's bid he should submit to maximize $\pi(v'|v, \hat{b}_n, n)$, given:

$$\left[v - \beta(v', \hat{b}_n, n) \right] \left[\Pi_{k=n+1}^N F^S(S'_k(v')) + \bar{F}_{n,v'}(\beta(v', \hat{b}_n, n)) \right] + \int_{\beta(v', \hat{b}_n, n)}^v (v - x) \bar{f}_{n,v'}(x) dx \quad (2)$$

where $\bar{F}_{n,v'}(t) = \Pi_{k=n+1}^N \int_0^t f^V(x) (1 - F^S(S'_k(v')|x)) dx$ is the probability that entry occurs and that the maximum value of all future entrants, when the incumbent's value at the end of round n is believed to be v' , is less than t , and $\bar{f}_{n,v'}(t) = \frac{\partial \bar{F}_{n,v'}(t)}{\partial v}$.¹⁸ The first part of equation 2 reflects the incumbent's profits when either there is no more entry or, all future entrants have values less than the jump bid, so the incumbent can win at his jump bid. The second part reflects the incumbent's profit if an entrant has a value above the jump bid. Differentiating equation 2 with respect to v' , and requiring that the first order condition equals zero when $v = v'$ (so that local incentive compatibility constraints are satisfied), gives

¹⁸Note that $\bar{F}_{n,v'}(\bar{V}) = 1 - \Pr[\text{no entry in future}]$, so that we are not double counting in equation 2.

the differential equation that defines the bid function:

$$\frac{d\beta(\cdot)}{dv} = \frac{[v - \beta(\cdot)] \left[\frac{d[\Pi_{k=n+1}^N F^S(S'_k(v))]}{dv} + \frac{\partial \bar{F}_{n,v}(\beta(\cdot))}{\partial v} \right] + \int_{\beta(\cdot)}^v (v - \hat{v}) \frac{\partial \bar{F}_{n,v}(\hat{v})}{\partial v} d\hat{v}}{\Pi_{k=n+1}^N F^S(S'_k(v)) + \bar{F}_{n,v}(\beta(\cdot))} \quad (3)$$

The lower boundary condition is provided by the condition that incumbents with values less than or equal to the standing bid will submit the standing bid. For values on $[\bar{V} - K, \bar{V}]$ bidders pool and submit $\beta(\bar{V} - K, \hat{b}_n, n)$.

When an incumbent bids $b \leq \beta(\bar{V} - K, \hat{b}_n, n)$, the posterior belief of any potential entrant in round $m > n$ about the incumbent's value will place all of the weight on $\beta^{-1}(b, \hat{b}_n, n)$. Thus, this potential entrant's entry threshold $S_m^{*'}(\beta^{-1}(b, \hat{b}_n, n))$ is implicitly defined by the following zero profit condition:

$$K = \int_{\beta^{-1}(b, \hat{b}_n, n)}^{\bar{V}} \pi(x|x, \beta^{-1}(b, \hat{b}_n, n), m) f^V(x|S_m^{*'}(\beta^{-1}(b, \hat{b}_n, n))) dx \quad (4)$$

Upon observing a bid at $\beta(\bar{V} - K, \hat{b}_n, n)$, beliefs will be consistent with Bayes Rule and a potential entrant will not participate.

Given the nature of this equilibrium we can solve the game recursively. For the final potential entrant, who believes that he will win if his value is greater than the incumbent's (in which case the final price will be the incumbent's value), we can solve for the equilibrium entry thresholds on a grid of possible values for an incumbent firm. Next we consider the previous potential entrant, and, given these final round thresholds, we can solve for both this entrant's entry thresholds and its equilibrium bid functions for a grid of possible values for an incumbent using equations 3 and 4. We then repeat the procedure for the third-from-last potential entrant, and so on until we reach the first round, where there will be no incumbent.

2.3.3 Equilibrium and Refinement

We now explain why the equilibrium just described exists and is the only equilibrium consistent with the D1 refinement (Banks and Sobel (1987), Cho and Kreps (1987)), which is a commonly used refinement for signaling models (Fudenberg and Tirole (1991)). For clarity we first consider the case with two potential entrants, which matches existing models in the signaling literature closely. We then consider the extension to the case with more firms. As the distribution of values has the same support for all types, adding more types has no effect on our arguments, so we assume there is only one type to reduce notation.

Our arguments will make use of three properties of the game. Let $\pi_v(b_1, S'_2)$ be the

expected profit of a first round incumbent with value v where b_1 is his bid and S'_2 is the entry threshold chosen by the second round potential entrant. Given our assumptions and the dominant strategies in the knockout game, $\pi_v(b_1, S'_2)$ will be continuous and differentiable in both arguments. The three properties are:

1. $\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} > 0$;
2. $\frac{\partial \pi_v(b_1, S'_2)}{\partial b_1} / \frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2}$ is monotonic in v ;
3. S'_2 is uniquely defined for any belief about the first potential entrant's value, and the potential entrant's response is more favorable to the incumbent when the potential entrant thinks that the incumbent's value is higher.

The appendix shows that these properties hold in our model.

The results in Mailath (1987) imply that there is an unique separating equilibrium on the $[0, \bar{V} - K]$ interval which can be found using the differential equation given by equation 3 and the boundary condition in a continuous type signaling model when the single crossing condition (property 2 above) holds. His results do not rule out the possibility of pooling equilibria on this interval. However, Ramey (1996) (who extends the results in Cho and Sobel (1990) to the case of an unbounded action space and a continuum of types on an interval) shows that these three properties imply that only a separating equilibrium will satisfy the D1 refinement, so our equilibrium must be the only one satisfying D1. As noted by Mailath (1987), this equilibrium will also be the separating equilibrium which is least costly to the first round potential entrant.¹⁹

The conditions also imply that, if an incumbent with value $\bar{V} - K$ prefers $\beta(\bar{V} - K)$, which will stop all future entry, to a lower bid then all incumbents with values above $\bar{V} - K$ will prefer $\beta(\bar{V} - K)$ to lower bids. But, firms with values above $\bar{V} - K$ will also strictly prefer to bid $\beta(\bar{V} - K)$ than any higher bid (for any beliefs of the potential entrant following a higher bid), because by bidding $\beta(\bar{V} - K)$ the incumbent can get the asset for sure at a lower price. Therefore, in equilibrium all entrants with values greater than $\bar{V} - K$ pool.

Three (or More) Rounds We now consider a model with three potential entrants (arguments for more rounds would follow directly from this case). We make the natural simplification by restricting ourselves to equilibria where all potential entrants make the same inferences from a bid by an incumbent and incumbents only make jump bids in the first round

¹⁹Given that we show below that the sequential mechanism tends to generate higher revenues than the auction, the fact that we focus on the *least cost separating* equilibrium implies that our results may be conservative since other equilibria in the sequential mechanism would give even higher revenues to the seller.

that they enter.²⁰ The two period equilibrium discussed above would define strategies for the final two rounds if the second period entrant enters and defeats any incumbent entrant from the first round (with an adjusted boundary condition to reflect the new standing bid). It therefore only remains to be shown that there is a unique sequential equilibrium bid function, which is fully separating for values $[0, \bar{V} - K]$, for a first round entrant. A first round entrant's jump bid sends a signal to the second round potential entrant, and, if he is still an incumbent in the final round, which must be the case if he is to win, the final potential entrant. Conditional on the incumbent surviving the second round, the third round is just a repeat of another two round game. The first round entrant's expected profit function is now $\pi_v(b_1, S'_2, S'_3)$, and the following properties hold:

1. $\frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial S'_2} > 0$ and $\frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial S'_3} > 0$;
2. $\frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial b_1} / \frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial S'_2}$ and $\frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial b_1} / \frac{\partial \pi_v(b_1, S'_2, S'_3)}{\partial S'_3}$ are both monotonic in v ;
3. both S'_2 and S'_3 are uniquely defined for any belief about the first entrant's value, and both potential entrants' responses are more favorable to the first entrant when they think that its value is higher.

These conditions allow us to apply the D1 refinement to the signaling game between the incumbent making the jump bid and every subsequent potential entrant to identify the unique separating equilibrium.

2.3.4 Illustrative Example of the Sequential Mechanism's Equilibrium

To provide additional clarity about how the mechanism works, given equilibrium strategies, Table 1 presents what happens in a game with four potential entrants and one type of firm with values distributed proportional to $LN(4.5, 0.2)$ on $[0, 200]$, $K = 1$ and $\sigma_\varepsilon = 0.2$ ($\alpha = 0.5$).

In the example, the first potential entrant enters if he receives a signal greater than 75.0, which is the case here. The signal thresholds in later rounds depend on the number of rounds remaining and the incumbent's value. So, when the incumbent is the same as in the previous round, the threshold S'^* falls since the expected profits of an entrant who beats the incumbent rise (because he will face less competition in the future). On the other hand, S'^* does not depend on the level of the standing bid given the incumbent's value, because it has no effect on the entrant's profits if he beats the incumbent in a knockout (since the standing bid must

²⁰It is possible that future potential entrants could ignore the information that they have on games before the last round. In this case, incumbents would choose to submit jump bids every round to signal information to the next potential entrant. This simplification allows us to consider a model where a firm sends at most one signal to many possible receivers.

Round	Initial	Potential Entrant				Post-Knockout	Post-Jump Bid
	Standing Bid	Value	Signal	S'^*	Entry	Standing Bid	Standing Bid
1	-	80.0	90.1	75.0	Yes	-	69.3
2	69.3	75.4	50.5	69.4	No	69.3	69.3
3	69.3	116.0	114.9	61.7	Yes	80.0	87.1
4	87.1	100.0	114.0	107.0	Yes	100.0	100.0
Seller's Revenue = 100.0, social surplus (winner's value less total entry costs) =113.0							

Table 1: A simple example of how the sequential mechanism works in a game with four potential entrants and one type of firm with values distributed proportional to $LN(4.5, 0.2)$ on $[0, 200]$, $K = 1$ and $\sigma_\varepsilon = 0.2$.

be below the incumbent's value). In round 2, the incumbent does not face entry, so there is no change in the standing bid because incumbents do not place additional jump bids.²¹ In round 3, the standing bid rises during the knockout, and the *new* incumbent places an additional jump bid. In round 4 the last potential entrant participates, but his value is less than the incumbent's and so revenue is the price at which this last entrant drops out.

We can also use this example to give intuition for how introducing selection affects bid functions and entry probabilities. With selection, the level of bids is determined by the fact that bids must be sufficiently high that firms with lower values will not want to copy them. In particular, if the entry decisions of later potential entrants are likely to be more sensitive to beliefs about the incumbent's value, then the equilibrium bid function must increase more quickly in v . A straightforward way to illustrate this is to focus on the last two rounds of the sequential mechanism when a new incumbent in the penultimate round only needs to worry about one more potential entrant, and the final round potential entrant would face no further entry if he enters and outbids the incumbent. This is illustrated in Figure 1, which compares the equilibrium bid functions in the penultimate round, and equilibrium probabilities of entry in the final round of the sequential mechanism for varying degrees of selection.

Specifically, the left panel displays bid functions for a new incumbent in the penultimate round, when the previous incumbent's value was 80. The right panel gives the probability that the final round potential entrant participates as a function of this new incumbent's value. Successively lower degrees of selection change the bid function so that when $\alpha \rightarrow 1$ it approaches the bid function in the LS (no selection) model (the bold line), which is a step function with a jump at a value of 119 (the level of the incumbent's value that deters all future entry). The slope of the bid function is more gradual for lower α s since the probability that the final round potential entrant participates declines more smoothly when α is lower.

²¹There would also have been no change in the standing bid if the entrant had come in, because the entrant's value was below the current bid, so the standing bid would not have risen in the knockout.

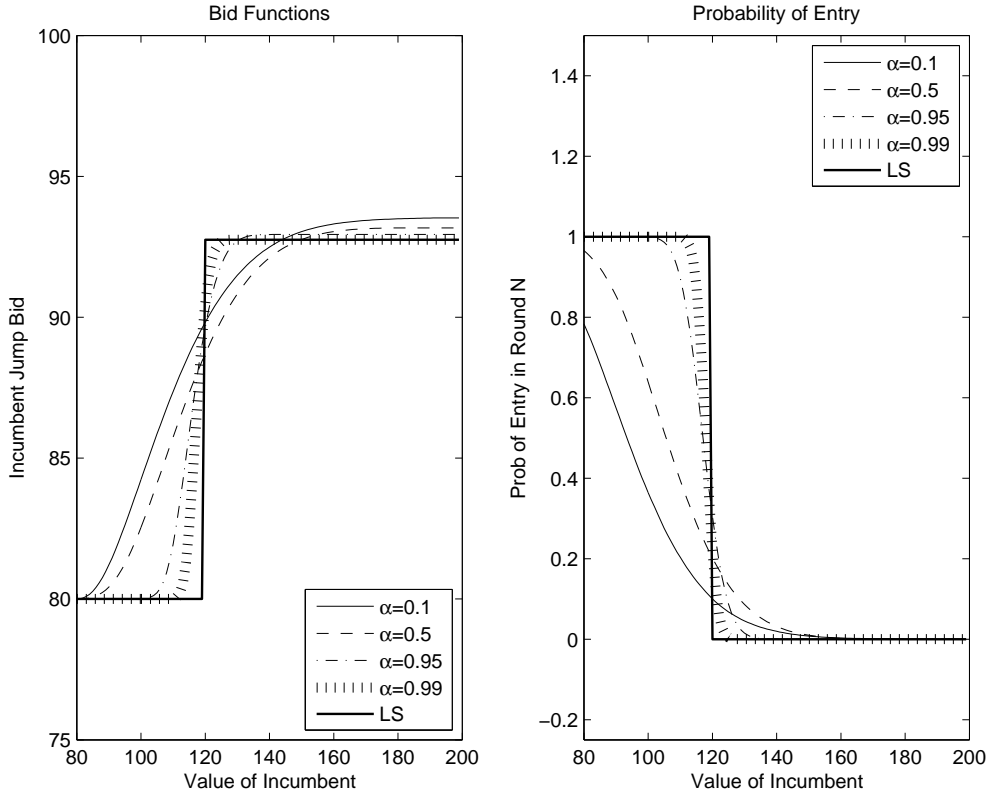


Figure 1: Penultimate round bid function for a new incumbent and probability of entry for the final round potential entrant, with symmetric firms, values $LN(4.5,0.2)$ on $[0,200]$, $K=1$ and standing bid of 80.

3 Comparison of Expected Revenues and Efficiency

Before introducing specific parameters estimated from data for USFS timber auctions, we present a more general comparison of expected revenues and efficiency between the sequential mechanism and the simultaneous entry auction. We see this general comparison as valuable, because it shows that our results in the empirical application are not going to be particularly sensitive to the parameters that we estimate, and they provide guidance about when auctions should perform well in other settings. Additionally, we allow a reserve price to be used in the simultaneous auction but restrict attention, for now, to a sequential mechanism with no reserve. In this way the results are biased *against* a seller preferring the sequential mechanism.

We focus on how the performance of the mechanisms depends on the level of entry costs (K) and the precision of the signal. We measure the precision of the signal by the parameter $\alpha = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_V^2}$, where a higher value of α indicates that signals are less precise. As a base case, we consider 8 symmetric firms whose values are distributed $LN(4.5,0.2)$ so that the value distribution has a mean of 91.6 and a standard deviation of 18.6.

Figure 2 shows the results of comparing expected revenues from the sequential mechanism

(with no reserve) and a simultaneous entry auction with an optimal reserve in (K, α) space.²² Filled squares represent outcomes where the expected revenues from the sequential mechanism are higher by more than 4% (of auction revenues), while hollow squares are outcomes where they are higher but only by between 0.1% and 4%. Diamonds represent cases where the simultaneous auction gives higher revenues. Crosses on the grid mark locations where the difference in revenues is less than 0.1%. Due to small numerical errors in solving differential equations and simulation error in calculating expected revenue, we take the conservative approach of not signing revenue differences in these cases.

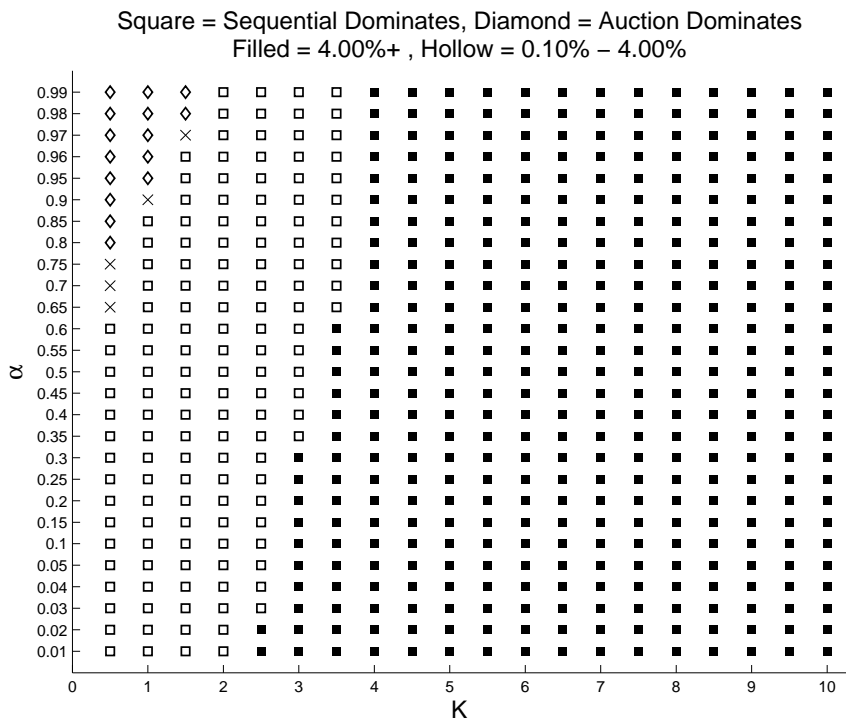


Figure 2: Expected revenue comparison for 8 symmetric firms, values LN(4.5,0.2), optimal reserve price in auction, no reserve price in sequential mechanism.

The results indicate that the sequential mechanism generally produces higher expected revenues than the auction, even when no reserve price is used in the sequential mechanism, whereas the optimal reserve price is used in the auction. The exception is for very low values of K and high α , but the revenue advantage of the auction is always small (the maximum difference is 1.1%). These points are consistent with BK's theoretical results as their model assumes no signals and requires that at least two firms will enter the auction, which implies that K must be low.

²²Sequential (auction) mechanism's expected revenues are calculated using 200,000 (5,000,000) simulations.

Economists and mechanism designers may be concerned with efficiency as well as revenues, and the sequential mechanism outperforms the auction along both dimensions. With no selection, low entry costs and symmetric bidders, BK show that the sequential mechanism is always more efficient, where efficiency is measured by the expected value of the winner less total entry costs paid. Increasing selection reduces the overall amount of entry and further wasteful entry costs, which serves to raise efficiency. It can also be the case that the expected value of winner in the sequential mechanism is higher than that in the auction, despite the lower number of entrants. For example, taking the parameters from Figure 2, when $K = 1$, the expected value of the winner is greater in the sequential mechanism for $\alpha \leq 0.25$. For higher K the expected value of the winner is always greater in the sequential mechanism. This is contrary to results from a no selection model, where “the expected value of the top bidder in the auction must be higher than in the sequential mechanism” (BK p. 1546). However, even when the expected value of the winner in the auction exceeds that in the sequential mechanism, the elimination of wasteful entry costs tends to sufficiently compensate so as to raise efficiency in the sequential mechanism. For example, over the grid given in Figure 2, there is no case in which the auction is more efficient than the sequential mechanism.

Regarding asymmetries, in a simultaneous entry auction weaker bidders need to consider the odds of competing against stronger bidders. While this is still true in the sequential mechanism, if the weaker bidders are approached last, given the separating equilibrium they know the values of the higher types that have entered. This permits more efficient entry of the weaker bidders and achieves a more efficient allocation of the good relative to the auction. For example suppose that $K = 5$ and $\alpha = 0.4$, $N = 4$ and the first two bidders approached have values proportional to $LN(4.5, 0.2)$, while the last two bidders approached have values proportional to $LN(4.4, 0.2)$.²³ The probability that each of the weaker firms enters the simultaneous auction is 0.20 and the probability that one of them wins is only 0.17. On the other hand, in the sequential mechanism the entry probabilities are 0.143 for the first one and 0.139 for the second and the probability that one of them wins is 0.28. This is much closer to the probability that one of the weaker firms will have the highest value (0.33). In this case, the sequential mechanism’s expected revenues of 83.34 exceed those of the auction, which are 78.40.

When bidders are asymmetric, sellers may prefer a first price auction with type-specific reserve prices to a second price auction with a uniform reserve. However, continuing with this example, even a first price auction with type-specific optimal reserve prices only generates expected revenues of 80.41, and so it is outperformed by the sequential mechanism with no

²³Our simulations show that approaching all of the high value firms first, followed by all of the low values firms is better than doing the opposite.

reserve price.

The separating equilibrium also has the effect of improving sellers' abilities to appropriate the larger rents in the sequential mechanism. This is important to note since, in BK's partial pooling equilibrium, high value participants' ability to completely forestall entry permits them to capture enough of the sequential mechanism's greater overall rents that sellers expect lower revenues than in the auction. Revenues in the auction are obviously determined by the second highest value of the bidders. In the sequential mechanism, revenues are determined by the maximum of the second highest value of participants and the deterring bid of the eventual winner. Thus, the larger rents can be appropriated either through stronger actual competition (the value of the second highest bidder is greater in the sequential mechanism) or by forcing the eventual winner to bid more aggressively to deter future potential competitors from participating. There are forces working towards and against encouraging stronger actual competition in the sequential mechanism. On the one hand, it can do better at selecting high value potential entrants into the mechanism. This is evidenced by the expected value of the winner in the sequential process sometimes exceeding that in the auction despite reduced entry. On the other hand, fewer bidders tend to participate, which lowers the expected value of the second highest value bidder. However, the threat of future potential entry by strong bidders forces the eventual winner to bid more aggressively and this tends to raise the seller's revenues. For example, the expected value of the second highest entrant is greater in the auction than in the sequential mechanism for each grid point in Figure 2. However, the winner's deterring bid is high enough to earn the seller higher revenues for almost all cases. In this sense, we find that the threat of future potential competition in the sequential mechanism leads to higher prices than does the greater actual competition in a simultaneous bid auction.

By way of example, Figure 1 illustrates how selection's effect on bid functions and equilibrium probabilities of entry serve to increase revenues in the sequential mechanism. Compared to the LS model, the bid functions with selection are higher for values less than 119, which is where this incumbent's value distribution is most dense. For example, when $\alpha = 0.1$, the mean and 90th percentile of a new incumbent who would find himself in the position of submitting a jump bid in the penultimate round are 101 and 121, respectively. While for a portion of the value distribution the bid functions are lower than when $\alpha = 1$, they again (slightly) exceed this bid function for very high values. This is because there is always a chance that the final potential entrant receives an optimistic signal and enters (unless it is inferred that the new incumbent's value is greater than $\bar{V} - K$, here 199), which cannot happen when the incumbent has a high value in BK's model. This is clearly illustrated in the right panel, which gives the probability that the final round potential entrant participates as

a function of the new incumbent’s value. This probability of entry is also a step function in the LS model. Once there is some selection, there is always a chance that the final round potential entrant participates, even if the incumbent is thought to have a high value (again, assuming its value is less than 199).

3.1 Sequential Mechanism with Reserve Prices

The numerical examples above indicate that a simple, stylized version of real-world sequential mechanisms tends to outperform the commonly used auction, even when the optimal reserve price is set in the auction. The sequential mechanism’s advantage over the auction could be increased through additional design elements, an obvious option being a reserve price. Figure 3 computes expected revenues when an optimal reserve price is added to each mechanism when there are five or eight symmetric bidders using the same value distribution parameters as before and assuming $K = 5$. For the sequential mechanism, only one reserve price is used, which is constant across all rounds in the mechanism. Generally, the seller could do better with a round-specific reserve price, but we view a constant reserve price as approximately imposing the same informational demands on the seller as does setting the optimal reserve price in the simultaneous auction.

Figure 3 shows that adding a constant reserve price to the sequential mechanism may substantially improve revenues. The reserve price affects sequential mechanism revenues in two ways. First, in the event that no firm has entered through the first $N - 1$ rounds, a reserve price guards the seller against giving the good away for free to the last potential entrant. Second, a reserve price raises the first entrant’s deterring bid function.

The effect of a reserve price varies across mechanisms and for different values of N and α . There are two main reasons for this. First, when entry is endogenous, a reserve price has a smaller impact when the level of entry is greater, as is generally the case (i) in the auction or (ii) when N is greater, as is clearly shown in Figure 3. Second, a reserve price excludes some bidders and if these were valuable to the seller, this reduces the value of a reserve price. This effect can be seen by noticing that the impact of a reserve price in the sequential mechanism falls for higher values of α : less selection implies that marginal and inframarginal entrants are more similar, which makes excluded bidders more valuable to the seller (it is also true that the level of entry increases in α , which also limits a reserve price’s impact).

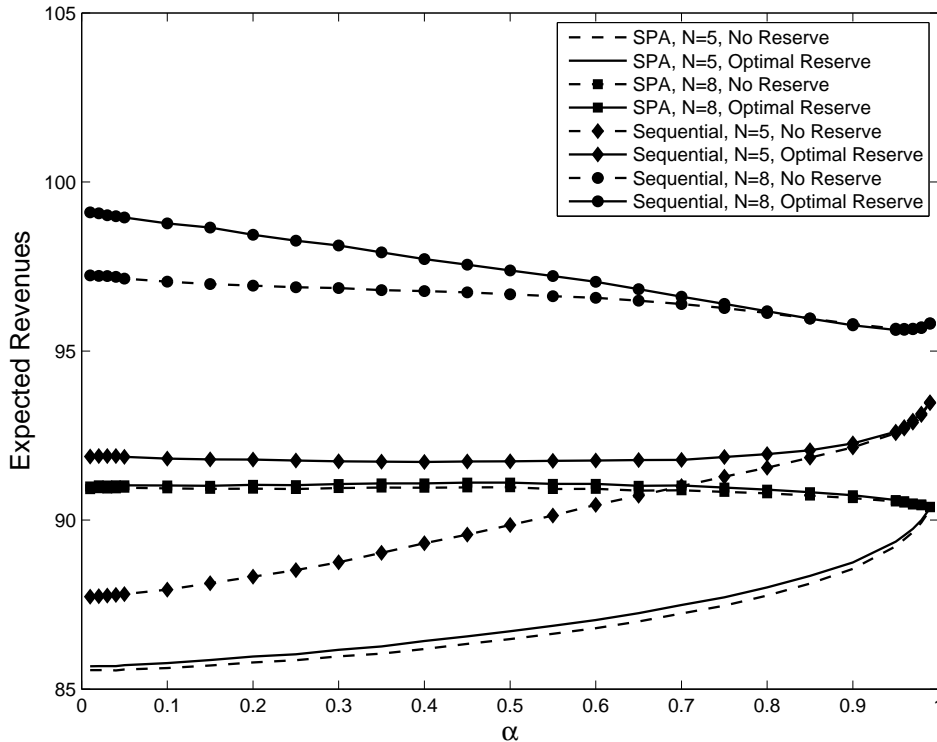


Figure 3: Expected revenue comparison for varying N , with and without reserves. Firms are symmetric with values distributed $\text{LN}(4.5, 0.2)$ and $K = 5$.

4 Empirical Application

We now turn to our empirical application that focuses on USFS timber auctions, which we view as a sensible environment for evaluating the effects of switching to a sequential mechanism. First, we find that these auctions are characterized by a costly and moderately selective entry process, features that we view as holding more generally across a wide variety of auction environments. Second, unlike other environments, such as the M&A market, we are able to observe the sale of many similar objects which facilitates estimation of bidder values. Third, while a great deal of work has concentrated on auction design tools, such as reserve prices, as means to increasing revenues in timber auctions, we show that a shift to a sequential sales process has a much larger impact. We are brief in our discussions of some reduced form evidence of selection and estimation method since Roberts and Sweeting (2011) provide a more detailed discussion of these topics. Additionally, Gentry and Li (2010) explore conditions under which an imperfectly selective entry model is non-parametrically identified for first price auctions.

4.1 Data

We analyze federal auctions of timberland in California. In these auctions the USFS sells logging contracts to individual bidders who may or may not have manufacturing capabilities (mills and loggers, respectively). When the sale is announced, the USFS provides its own “cruise” estimate of the volume and value of timber for each species on the tract as well as estimated costs of removing and processing the timber. It also announces a reserve price and bidders must indicate a willingness to pay at least this amount to qualify for the auction. After the sale is announced, interested potential bidders perform their own private cruises in order to assess the tract’s value. These cruises are informative about the tract’s volume, species make-up and timber quality.

We assume that bidders have independent private values. This assumption is also made in other work with similar timber auction data (see for example Baldwin, Marshall, and Richard (1997), Haile (2001) or Athey, Levin, and Seira (forthcoming)). A bidder’s private value is primarily related to its own contracts to sell the harvest, inventories and private costs of harvesting. In addition, we focus on the period 1982-1989 when resale, which can introduce a common value element, was limited (see Haile (2001) for an analysis of timber auctions with resale).

We also assume non-collusive bidder behavior. While there has been some evidence of bidder collusion in open outcry timber auctions, Athey, Levin, and Seira (forthcoming) find strong evidence of competitive bidding in these California auctions.

Our model assumes that bidders receive an imperfect signal of their value and they must pay a participation cost to enter the auction.²⁴ We interpret the USFS’s publicly available tract appraisal and a firm’s own knowledge of its sales contracts and capabilities as generating its pre-entry signal. Participation in these auctions is costly for numerous reasons. In addition to the cost of attending the auction, a large fraction of a bidder’s entry cost is its private cruise. People in the industry tell us that firms do not bid without doing their own cruise, which can provide information that bidders find useful, such as trunk diameters, but is not provided in USFS appraisals.

We use data on 887 ascending auctions.²⁵ Table 2 shows summary statistics for our sample. Bids are given in \$ per thousand board feet (mbf) in 1983 dollars. The average mill bid is 20.3% higher than the average logger bid. As suggested in Athey, Levin, and Seira (forthcoming), mills may be willing to bid more than loggers due to cost differences or the imperfect competition loggers face when selling felled timber to mills.

²⁴We note that we are not the first to model a costly entry decision into these auctions (e.g. Athey, Levin, and Seira (forthcoming)).

²⁵Roberts and Sweeting (2011) include a detailed description of the sample selection process.

Variable	Mean	Std. Dev.	25 th -tile	50 th -tile	75 th -tile	N
WINNING BID (\$/mbf)	86.01	62.12	38.74	69.36	119.11	847
BID (\$/mbf)	74.96	57.68	30.46	58.46	105.01	3426
LOGGER	65.16	52.65	26.49	49.93	90.93	876
MILL	78.36	58.94	32.84	61.67	110.91	2550
LOGGER WINS	0.15	0.36	0	0	0	887
FAIL	0.05	0.21	0	0	0	887
ENTRANTS	3.86	2.35	2	4	5	887
LOGGERS	0.99	1.17	0	1	1	887
MILLS	2.87	1.85	1	3	4	887
POTENTIAL ENTRANTS	8.93	5.13	5	8	13	887
LOGGER	4.60	3.72	2	4	7	887
MILL	4.34	2.57	2	4	6	887
SPECIES HHI	0.54	0.22	0.35	0.50	0.71	887
DENSITY (hundred mbf/acre)	0.21	0.21	0.07	0.15	0.27	887
VOLUME (hundred mbf)	76.26	43.97	43.60	70.01	103.40	887
RESERVE (\$/mbf)	37.47	29.51	16.81	27.77	48.98	887
SELL VALUE (\$/mbf)	295.52	47.86	260.67	292.87	325.40	887
LOG COSTS (\$/mbf)	118.57	29.19	99.57	113.84	133.77	887
MFCT COSTS (\$/mbf)	136.88	14.02	127.33	136.14	145.73	887

Table 2: Summary statistics for sample of California ascending auctions from 1982-1989. All monetary figures in 1983 dollars. SPECIES HHI is the Herfindahl index for wood species concentration. SELL VALUE, LOG COSTS and MFCT COSTS are USFS estimates of the value of the tract and the logging and manufacturing costs of the tract, respectively.

We define potential entrants as the auction’s bidders plus those firms who bid within 50 km of an auction over the next month. One way of assessing the appropriateness of this definition is that 98% of the bidders in any auction also bid in another auction within 50 km of this auction over the next month and so we are unlikely to be missing many actual potential entrants. The median number of potential bidders is eight (mean of 8.93) and this is evenly divided between mills and loggers.

In Table 2, entrants are defined as the set of bidders we observe at the auction, even if they did not submit a bid above the reserve price.²⁶ The median number of mill and logger entrants are three and one, respectively. Among the set of potential logger entrants, on average 21.5% enter, whereas on average 66.1% of potential mill entrants enter. The differences in bids and entry decisions are consistent with mills having significantly higher values than loggers.²⁷

²⁶However, in our preferred empirical specification below, we interpret the data more cautiously and allow bidders that do not submit bids to have entered (paid K), but learned that their value was less than the reserve price.

²⁷Roberts and Sweeting (2011) present evidence that differences in values, and not entry costs, explain

4.2 Evidence of Selection

Roberts and Sweeting (2011) present reduced form evidence that the data are best explained by a model allows for selection. There are two main pieces of evidence. First, Athey, Levin, and Seira (forthcoming) show that in the type-symmetric mixed strategy equilibrium of a model with endogenous, but non-selective, entry and asymmetric bidder types, whenever the weaker type enters with positive probability, the stronger type enters with probability one. Thus, for any auction with some logger entry, a model with no selection would imply that all potential mill entrants enter. In 54.5% of auctions in which loggers participate, and there are some potential mill entrants, some, but not all, mills participate. Likewise, they show that whenever the stronger type enters with probability less than one, a model with no selection implies that weaker types enter with probability zero. However, in the data we find that in 61.1% of auctions in which only some mill potential entrants participate and potential logger entrants exist, some loggers enter. A model with selective entry can rationalize partial entry of both bidder types into the same auction.

Second, a model without selection implies that bidders are a random sample of potential entrants. Roberts and Sweeting (2011) test this by estimating a Heckman selection model with the exclusion restriction that potential competition affects a bidder’s decision to enter an auction, but has no direct effect on values. The second stage regression of all bids on auction covariates and the estimated inverse Mills ratio from a first stage probit of the decision to participate shows a positive and significant coefficient on the inverse Mills ratio. This is consistent with bidders being a selected sample of potential entrants.

The evidence presented in this section strongly suggests that the entry process is selective. However, it does not pin down the degree of selection. Therefore, we now describe how we estimate our model to measure the degree of selection.

4.3 Estimation Using Importance Sampling

To take the model to data, we need to specify how the parameters of the model may vary across auctions, as a function of observed auction characteristics and unobserved heterogeneity. Both types of heterogeneity are likely to be important as the tracts we use differ greatly in observed characteristics, such as sale value, size and wood type, and they also come from different forests over several years so they may differ in other characteristics as well. Both observed and unobserved (to the econometrician) heterogeneity may affect entry costs and the degree of selection, as well as mean values.²⁸

why mills are more likely than loggers to enter an auction.

²⁸In the specification below, α is not a function of observables as when we allowed for this the estimated effects of observables on the degree of selection were small and imprecise.

Our estimation approach is based on Akerberg (2009)’s method of simulated maximum likelihood with importance sampling. We fully describe our estimation method in the appendix and in Roberts and Sweeting (2011). However, here we note several features of our specification.

We assume that the parameters are distributed across auctions according to the following distributions, where X_a is a vector of observed auction characteristics and $TRN(\mu, \sigma^2, a, b)$ is a truncated normal distribution with parameters μ and σ^2 , and upper and lower truncation points a and b .

$$\text{Location Parameter of Logger Value Distribution: } \mu_{a,\text{logger}} \sim N(X_a\beta_1, \omega_{\mu,\text{logger}}^2)$$

$$\text{Difference in Mill/Logger Location Parameters: } \mu_{a,\text{mill}} - \mu_{a,\text{logger}} \sim TRN(X_a\beta_3, \omega_{\mu,\text{diff}}^2, 0, \infty)$$

$$\text{Scale Parameter of Mill and Logger Value Distributions: } \sigma_{V_a} \sim TRN(X_a\beta_2, \omega_{\sigma_V}^2, 0.01, \infty)$$

$$\alpha: \alpha_a \sim TRN(\beta_4, \omega_{\alpha}^2, 0, 1)$$

$$\text{Entry Costs: } K_a \sim TRN(X_a\beta_5, \omega_K^2, 0, \infty)$$

These specifications reflect our assumptions that σ_V , α and K are the same for mills and loggers within any particular auction, even though they may differ across auctions.

To apply the estimator, we also need to define the likelihood function based on the open outcry auction data. Two problems arise when interpreting these data. First, a bidder’s highest announced bid in an open outcry auction may be below its value, and it is not obvious which mechanism leads to the bids that are announced (Haile and Tamer (2003)). Second, if a firm does not know its value when taking the entry decision, it may learn (after paying the entry cost) that its value is less than the reserve price and so not submit a bid. We take a conservative approach (the details of which are provided in the appendix) when interpreting the data by assuming that the winning bidder has a value greater than the second highest bid, the second highest observed bid is equal to the value of the second-highest bidder, all other bidders had values less than the highest observed bid and that potential entrants that we do not see bid may or may not have paid the entry cost.

5 Empirical Results

In this section we present estimates of our structural model and counterfactual results measuring the benefits to the USFS of switching from the current simultaneous entry and simultaneous bid auction to our simple sequential process.

5.1 Parameter Estimates

Table 3 presents the parameter estimates for our structural model.²⁹ We allow the USFS estimate of sale value and its estimate of logging costs to affect mill and logger values and entry costs since these are consistently the most significant variables in regressions of reserve prices or winning bids on observables, including controls for potential entry. We also control for species concentration since our discussions with industry experts lead us to believe that can matter to firms. We allow for auction-level unobserved heterogeneity (to the econometrician) in all parameters. The righthand columns show the mean and median values of the structural parameters when we take 10 simulated draws of the parameters for each auction. For the rest of the paper, we refer to these as the “mean” and “median” values of the parameters. All standard errors are based on a non-parametric bootstrap, where both auctions and draws are re-sampled, with 100 repetitions.

The coefficients show that tracts with greater sale values and lower costs are more valuable, as one would expect. There is unobserved heterogeneity in values across auctions (the standard deviation of μ_{logger}) and some unobserved heterogeneity in the difference between mill and logger mean values across auctions (the standard deviation of $\mu_{\text{mill}} - \mu_{\text{logger}}$).

Based on the mean value of the parameters, the mean values of mills and loggers in the population are, in 1983 dollars, \$61.95/mbf and \$42.45/mbf, respectively, a 46% difference. We estimate a mean entry cost of \$2.05/mbf, also in 1983 dollars. One forester we spoke with estimated modern day cruising costs of approximately \$6.50/mbf, or \$2.97/mbf in 1983 dollars. It is sensible that our estimate is less than the forester’s estimate if firms in our data are able to use any information they learn when deciding whether to enter other auctions.

Our estimates of the α s across auctions indicate a moderate amount of selection in the data. This is illustrated by the difference in expected values for marginal and inframarginal bidders in a representative auction where the reserve price and the number of potential mill and logger entrants are set to their respective medians of \$27.77/mbf, four and four. Based on the mean parameter values, the expected values of a marginal and inframarginal mill entrant are \$45.22/mbf and \$68.13/mbf, respectively (the former is lower than the population average because most mills enter). The comparable numbers for loggers are \$48.13/mbf and \$59.80/mbf, respectively.

Our estimation approach assumes that, if there are multiple equilibria, the firms will play the equilibrium where mills have the lower S^* . We can check whether our parameter estimates can support multiple equilibria by plotting type-symmetric “equilibrium best response functions” for mills and loggers for each auction. For each auction, our parameter estimates

²⁹Roberts and Sweeting (2011) discuss alternative estimation methods that were attempted, such as Nested Pseudo-Likelihood, and model fit.

Parameter	β parameters				ω parameter		
	Constant	log SELL VALUE	log LOG COSTS	SPECIES HHI		Mean	Median
$\mu_{a,\text{logger}}$ $\sim N(X_a\beta_1, \omega_{\mu,\text{logger}}^2)$	-9.6936 (1.3690)	3.3925 (0.1911)	-1.2904 (0.1332)	0.2675 (0.1386)	0.3107 (0.0213)	3.5824 (0.0423)	3.5375 (0.0456)
$\mu_{a,\text{mill}} - \mu_{a,\text{logger}}$ $\sim TRN(X_a\beta_3, \omega_{\mu,\text{diff}}^2, 0, \infty)$	3.6637 (0.8890)	-0.4998 (0.1339)	-0.0745 (0.0919)	-0.1827 (0.1007)	0.1255 (0.0163)	0.3783 (0.0242)	0.3755 (0.0249)
σ_{V_a} $\sim TRN(X_a\beta_2, \omega_{\sigma_v}^2, 0.01, \infty)$	4.0546 (0.7872)	-0.7379 (0.0994)	0.1393 (0.1025)	0.0895 (0.0813)	0.0796 (0.0188)	0.5763 (0.0273)	0.5770 (0.0302)
α_a $\sim TRN(\beta_4, \omega_{\alpha}^2, 0, 1)$	0.7127 (0.0509)	-	-	-	0.1837 (0.0446)	0.6890 (0.0362)	0.6992 (0.0381)
K_a $\sim TRN(X_a\beta_5, \omega_K^2, 0, \infty)$	1.9622 (13.2526)	-3.3006 (2.7167)	3.5172 (2.4808)	-1.1876 (1.5721)	2.8354 (0.6865)	2.0543 (0.2817)	1.6750 (0.3277)

Table 3: Simulated maximum likelihood with importance sampling estimates allowing for non-entrants to have paid the entry cost. The rightmost columns show the mean and median values of the structural parameters when we take 10 simulated draws of the parameter for each auction. Standard errors based non-parametric bootstrap with 100 repetitions. $TRN(\mu, \sigma^2, a, b)$ is a truncated normal distribution with parameters μ and σ^2 , and upper and lower truncation points a and b . Based on 887 auctions.

support only a single equilibrium. This is because our estimates imply a large difference in the mean values of loggers and mills, relatively low entry costs and a moderate amount of selection, all of which promote a unique equilibrium.

5.2 Counterfactual Results

Table 4 compares expected revenues and efficiency from the sequential mechanism and the simultaneous entry auction for a range of parameters and different numbers of firms. The simulations assume mills are approached first (in a random order) followed by loggers, although we have found some cases where a different order can strengthen the results below.

The first line in Table 4 gives the results for the representative auction (four mills and four loggers) based on the mean parameter estimates used in the constructing the figures above. Relative to setting no reserve price in the simultaneous entry, simultaneous bid auction, the sequential mechanism with no reserve price improves the USFS's revenues by 1.81%. For a tract of average size (7,626 mbf) the expected revenue difference would be \$9,834.

The increase in revenues in this representative case of switching from the simultaneous bid auction with no reserve price to the sequential mechanism with no reserve price is 9.05 times as large as the improvement from using an optimal reserve in the simultaneous bid auction, which is just 0.2%. The finding that the revenue increase from using the sequential mechanism is much larger than the returns to using a reserve price in the current auction format is important since understanding optimal reserve price policies for timber auctions has been the subject of significant interest (examples include Mead, Schniepp, and Watson (1981), Paarsch (1997), Haile and Tamer (2003), Li and Perrigne (2003) and Aradillas-Lopez, Gandhi, and Quint (2010)). Additionally, the sequential mechanism provides an easily implementable mechanism that does not require the USFS to possess detailed information on all of the model's primitives. Such information would be required to set an optimal reserve price. However, were the USFS to possess such information, a reserve price could also be set in the sequential mechanism. If a reserve price is used in the sequential mechanism, the increase in revenues becomes 10.43 times as large as the gain to setting an optimal reserve price in the auction. This advantage would increase if we considered round-specific, optimal reserve prices in the sequential mechanism.

Not only does the sequential mechanism have a much larger impact on revenues than does setting an optimal reserve price in the standard auction format, it also increases efficiency, as shown in the penultimate column in Table 4. In the representative case given in the first row of the table, the USFS captures the majority of the increase in surplus, but expected firm profits still increase in the sequential mechanism. As mentioned in Section 3, the sequential mechanism tends to promote more efficient entry of weaker bidders and this increases their

Case	Parameters					Expected Revenues (\$/mbf)			Expected Efficiency			
	N_{mill}	N_{logger}	μ_{logger}	μ_{diff}	σ_V	α	K	SPA $R=0$	SEQ $R=0$	SPA $R=R^*$	Total Δ (\$/mbf)	% of $\Delta \rightarrow$ to Firms
	4	4	3.582	0.378	0.576	0.689	2.05	71.25	72.54	71.39	1.88	31.38
2	1	4	3.582	0.378	0.576	0.689	2.05	50.96	52.32	51.63	1.45	6.21
3	7	4	3.582	0.378	0.576	0.689	2.05	83.91	85.21	83.98	2.16	39.81
4	4	0	3.582	0.378	0.576	0.689	2.05	64.42	64.46	64.69	0.91	95.60
5	4	8	3.582	0.378	0.576	0.689	2.05	75.87	77.60	75.97	2.31	25.11
6	4	4	<i>2.921</i>	0.378	0.576	0.689	2.05	34.10	35.50	34.34	1.58	11.39
7	4	4	<i>4.243</i>	0.378	0.576	0.689	2.05	142.75	143.75	142.81	2.25	55.56
8	4	4	3.582	<i>0.169</i>	0.576	0.689	2.05	61.90	62.93	62.02	1.88	45.21
9	4	4	3.582	<i>0.587</i>	0.576	0.689	2.05	83.91	85.32	84.07	1.66	15.06
10	4	4	3.582	0.378	<i>0.349</i>	0.689	2.05	57.13	58.79	57.47	1.51	-9.93
11	4	4	3.582	0.378	<i>0.804</i>	0.689	2.05	91.08	91.64	91.16	1.98	71.72
12	4	4	3.582	0.378	0.576	<i>0.505</i>	2.05	70.64	72.36	71.02	1.95	11.79
13	4	4	3.582	0.378	0.576	<i>0.872</i>	2.05	71.89	73.08	71.91	2.03	41.38
14	4	4	3.582	0.378	0.576	0.689	<i>0.39</i>	75.74	75.71	75.75	0.54	105.56
15	4	4	3.582	0.378	0.576	0.689	<i>3.72</i>	66.69	69.32	67.12	2.93	10.24

Table 4: Comparing the impact of the sequential mechanism on expected revenues and efficiency. The first line shows the results for the representative auction with four mills and four loggers and the mean parameter estimates from Table 3. This is what we refer to as the “representative case”. Italics indicate changes from this representative auction and these changes are ± 1 standard deviation changes based on our estimates of the distributions of the structural parameters. SEQ refers to the sequential mechanism, SPA refers to the auction, $R=0$ and $R=R^*$ indicates that either no reserve price or the optimal reserve price is used, respectively. The first first three columns following the parameters give the expected revenues in the three mechanisms. The final two columns compare the expected efficiency of a sequential mechanism and an auction with no reserve. The first of these is the total change in efficiency. The second gives the percent of the change which accrues to firms. Auction results based on 5,000,000 simulations and sequential results based on 200,000 simulations.

expected profits. In the USFS auctions, switching from the current auction format to the sequential mechanism tends to increase expected logger profits without substantially harming those of mills. For example, in the representative case, expected logger profits increase 21% when the sequential mechanism is used, while mill profits only fall by 0.60%.

The other rows in the table compare outcomes when we increase or decrease the number of potential entrants or structural parameters by one standard deviation (the changing parameter is in *italics*), reflecting the fact that our estimates imply that the coefficients will differ across sales. The cases we consider indicate that using the sequential mechanism generally raises expected revenues. In case 14 the entry cost is very low and in either mechanism almost all firms participate so that revenues are essentially the same. In all cases, once a constant reserve price is used in the sequential mechanism, it earns higher revenues than the current auction format even with an optimal reserve price. We can see that setting a reserve price in the standard auction format is particularly ineffective when there are many potential entrants or when entry is less selective (α is high). In all cases the sequential mechanism increases efficiency and in only one example does total bidder surplus fall (case 10). The finding from the first row that loggers benefit from switching to the sequential mechanism holds in all rows. Additionally, when expected mill profit falls, it tends to be by a small amount, and in some cases it rises. As an example, in case 8, when μ_{diff} is low (0.169), loggers' expected profit increases by 10.18% and mills' increases by 1.40%.

The USFS also uses first price, sealed bid auctions to sell timber. We can also compare the performance of the sequential mechanism to this alternative. Across all of the cases in Table 4, with the exception of cases 6 and 14, a sequential mechanism with no reserve price earns the USFS higher revenues than a first price auction with an optimal reserve price. Introducing a reserve price to the sequential mechanism increases its advantage over the first price auction by even more so that it now dominates in all cases.

While we believe that the simultaneous entry auction is the natural way to think about how USFS auctions currently operate, we have also computed expected auction revenues if, instead, firms enter sequentially (in the same order as the sequential mechanism) before simultaneously submitting bids. For the representative auction, expected revenues in this case are \$71.84/mbf which is still less than the revenues from the sequential bidding mechanism. This pattern holds more generally in the other rows in Table 4 where we were able to solve a sequential entry auction game: in only two cases (6 and 14) did the sequential entry auction give higher revenues than the sequential mechanism with no reserve price and in both cases the differences were small (\$0.40/mbf and \$0.06/mbf, respectively).³⁰

³⁰As explained in footnote 14, there is a high computational burden to solving the sequential entry auction because it is necessary to solve for a threshold as a function of all possible histories of the game simultaneously. For cases 3, 5, 13 and 15 we could not do so satisfactorily.

Government sales and procurement programs often have distributional requirements that a certain portion of contracts be awarded to targeted firms. The US federal government seeks to award 23% of the \$400 billion worth of annual contracts to small businesses (see Athey, Coey, and Levin (2011) for additional discussion). The primary ways of favoring smaller businesses are through set-asides, where only targeted firms can participate, and bid subsidies for preferred firms. The USFS has historically used set-asides and recent work (Athey, Coey, and Levin (2011)) suggests this may come at a substantial revenue and efficiency loss relative to using bid subsidies. In this light, our findings that the USFS can increase revenues, efficiency and the profits of loggers with only small decreases in mill profits (and sometimes increases) by switching from the current auction format to a sequential mechanism may be particularly useful. Although a full comparison of bid subsidies, set-asides and the sequential mechanism is beyond the scope of this paper, we see our findings as suggesting that the sequential mechanism may present procurement agencies with an effective alternative method for allocating projects to targeted bidders. Additionally, the sequential mechanism requires only that the agency be able to identify targeted firms, which is also required in the use of set-asides and subsidies, and does not require determining optimal subsidy amounts or even setting reserve prices. For timber auctions, we have been told by USFS officials that they believe that they can accurately identify the set of potential entrants for any given sale. Even if at times they are unsure, it would be straightforward to allow potential participants to costlessly identify themselves before the full details of a sale are announced.

Our discussion so far has largely ignored potential practical impediments to implementing the sequential mechanism. First, were the USFS to use the sequential mechanism, there may be concern that approaching firms in an order places some of them at an advantage over others and may lead firms to try to affect the order in which they are approached. However, for all of the examples that we have considered, expected firm profits are fairly constant across the order of moves within bidder type, and there is no systematic pattern suggesting that a particular spot in the order is best. Intuitively, while the first potential entrant will be more likely to participate, he also must pay more to win. For example, in the representative auction, where the four mills are approached first followed by the four loggers, the expected profits (in \$/mbf) by order are {6.07, 6.09, 6.14, 6.18, 1.08, 1.05, 1.09, 1.04}. The maximum amount by which expected mill (logger) profits differ in this case is 0.016 (0.042). Second, there may be some concern about whether the USFS can commit to an order. However, repeated use of the mechanism likely would incentivize the USFS to maintain its credibility through consistent commitment to stated orders. Additionally, the lack of variation of profits across spots in the order could mean that firm lobbying efforts, which might dissuade a seller from sticking to a stated order, are likely to be small. Third, collusion may be a concern

given the existing evidence from other USFS regions consistent with noncompetitive bidding (Athey, Levin, and Seira (forthcoming)). However, as Bulow and Klemperer (2009) note (their footnote 40), the “simple auction is perhaps more easily undermined, than a sequential process, by collusion.”

There may also be some concern that switching to the sequential mechanism would greatly increase the time required to sell any stand of timber. While the length of the bidding process would necessarily increase, we note that there is already a sizable gap (over a month) between when a sale is announced and when it is completed.³¹ Since cruising takes between a day and seven days, depending on the size of the sale, even in the extreme (assuming a large sale in which 8 potential bidders all decide to participate), a sequential mechanism could be run in under two months. Often the sequential process could happen much faster, but even an extra month may be a small price to pay to realize the sequential mechanism’s advantages.

We have also ignored the USFS’s cost of switching to the sequential mechanism. Although the cost of selling a stand of timber is likely to be similar across mechanisms, there may be a fixed cost associated with switching from the currently used format to a sequential process, which would have to be measured against the potential gains from doing so. Based on the 15 cases in Table 4 alone, USFS revenues would increase by approximately \$315,000 (in 2011 dollars) compared to using the current auction format with no reserve price.³² Given that these 15 sales represent less than 0.3% of the tracts sold by the USFS in CA between 1982 and 1989, this one-time, sunk cost is likely to be small relative to the associated increase in revenues.

6 Conclusion

This paper compares the performance of a sequential and a simultaneous bidding mechanism in an environment where it is costly for potential buyers to participate and they receive imperfectly-informative signals about their values prior to deciding whether to enter, so that the entry process is selective. In contrast to results when there is no selection, a very simple sequential mechanism can generate higher expected revenues for the seller than the commonly used auction, and it also has an efficiency advantage so that buyers may prefer it is as well. The revenue result holds even though there is less entry (actual competition) into the sequential mechanism. Instead, with selection, the sequential mechanism can do a better job of allocating the good to the firm with the highest value and this fact, combined

³¹In fact the gap is usually much longer since the USFS must file documents to comply with the National Environmental Policy Act.

³²These calculations assume the average tract size in the 15 cases is 7,626 mbf.

with the feature that firms with high values have to bid aggressively in an attempt to deter future entry, provides its revenue advantage.

We view our results as relevant and important for at least three reasons. First, a selective entry process is likely an appropriate description of many real-world settings where a bidding process is used to sell an asset. This is because potential buyers often possess some pre-existing knowledge of their match for the asset, but will need to conduct costly additional research to determine how much they should be willing to pay. In these cases, our results point to conclusions about how bidding should be structured that are different to those in the existing literature. Our model also allows us to explain certain features of the data, such as jump-bids not deterring all future entry in takeover contests, so that multiple jump bids, sometimes by different firms, are observed (e.g., Betton and Eckbo (2000) or Betton, Eckbo, and Thorburn (2008)). These facts cannot be explained by a model with no selection.

Second, the revenue differences that we identify are not trivial. As a comparison, we consider the seller's return to setting an optimal reserve price in a simultaneous auction, which is the type of relative small design change that is the focus of the existing empirical literature. For the representative auction in our data, we estimate that the seller's return to switching to the sequential mechanism would be nine times greater than the return to setting the optimal reserve price. The absolute difference in revenues can also be large when entry costs are higher or entry is more selective than we estimate to be the case in USFS auctions.

Third, our results are directly relevant to an on-going legal debate about how corporate sales should be structured in order to allow boards to fulfill their *Revlon* duties to maximize shareholder value. At the very least, our results suggest that there are circumstances in which a sequential bidding process will achieve this more effectively than a simultaneous one, and they highlight two factors (entry costs and selection) on which the results are likely to depend. One concern that has been raised with sequential processes is that all potential buyers are not treated equally, so that firms that move first may be able to deter later ones and retain a right to match the prices offered by any later competition that emerges. This is true in our model, but it does not necessarily mean that the firms that move first earn higher revenues. In fact, our results suggest that expected payoffs are fairly equal across the order in the presence of selection, because early movers also pay entry costs more often and are less likely to win when they enter.

One might believe that while simultaneous auctions often operate in exactly the way modeled here, the stylized sequential mechanism that we consider is not widely implemented in its exact form, perhaps suggesting that it is impractical or has some hidden disadvantage. We do not believe this to be the case. The only thing that the seller needs to know is the set of potential buyers, and in many cases it would be straightforward for these firms to identify

themselves. The seller does need to be able to commit to approaching potential buyers in a particular order, and to develop a system for distributing information about previous bids. For many assets, any costs involved are likely to be small, and for a firm or government agency involved in repeated transactions (e.g., procurement) they would be spread over a large number of contracts. Instead, a more plausible reason for why the exact sequential mechanism considered here is not used is that there are alternative sequential mechanisms that can do even better, consistent with the fact that the seller optimal mechanism is almost certainly some sort of sequential search process and that, within the sequential mechanism, unlike the simultaneous auction, the ability to set a reserve price, and possibly other design elements, can increase seller revenues substantially.

There are, of course, some limitations of the model that we consider here, which may be important in some real-world settings. For example, our IPV assumption will not be satisfied for assets where potential buyers have to form imperfect opinions about some innate future potential. A common value component would change strategies significantly in the sequential mechanism as the incumbent bidder could signal that he believes the common value to be low in order to deter entry. We also assume that firms act competitively, while the structure of the selective mechanism might affect incentives for collusion on either entry decisions or bids. Understanding how these factors would affect the relative performance of sequential and simultaneous mechanisms appear to be profitable directions for future research.

References

- ACKERBERG, D. (2009): “A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation,” *Quantitative Marketing and Economics*, 7, 343–376.
- ARADILLAS-LOPEZ, A., A. GANDHI, AND D. QUINT (2010): “Identification and Testing in Ascending Auctions with Unobserved Heterogeneity,” *Working Paper, University of Wisconsin*.
- ATHEY, S., D. COEY, AND J. LEVIN (2011): “Set-Asides and Subsidies in Auctions,” *Working Paper*.
- ATHEY, S., P. CRAMTON, AND A. INGRAHAM (2003): “Upset Pricing in Auction Markets: An Overview,” White Paper, Market Design, Inc., on behalf of British Columbia Ministry of Forests.
- ATHEY, S., J. LEVIN, AND E. SEIRA (forthcoming): “Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions,” *Quarterly Journal of Economics*.
- BAJARI, P., H. HONG, AND S. RYAN (2010): “Identification and Estimation of a Discrete Game of Complete Information,” *Econometrica*, 78(5), 1529–1568.
- BALDWIN, L., R. MARSHALL, AND J.-F. RICHARD (1997): “Bidder Collusion at Forest Service Timber Sales,” *Journal of Political Economy*, 105(4), 657–699.
- BANKS, J., AND J. SOBEL (1987): “Equilibrium Selection in Signalling Games,” *Econometrica*, 55(647-662).
- BERRY, S. (1992): “Estimation of a Model of Entry in the Airline Industry,” *Econometrica*, 60(4), 889–917.
- BETTON, S., AND B. ECKBO (2000): “Toeholds, Bid Jumps, and Expected Payoffs in Takeovers,” *Review of Financial Studies*, 13, 841–882.
- BETTON, S., B. ECKBO, AND K. THORBURN (2008): “Merger Negotiations and the Toehold Puzzle,” *Journal of Financial Economics*, 91, 158–178.
- BOONE, A. L., AND J. H. MULHERIN (2007): “How Are Firms Sold?,” *Journal of Finance*, 62(2), 847–875.
- BULOW, J., AND P. KLEMPERER (2009): “Why Do Sellers (Usually) Prefer Auctions?,” *American Economic Review*, 99(4), 1544–1575.
- CHO, I., AND D. KREPS (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102(2), 179–222.
- CHO, I., AND J. SOBEL (1990): “Strategic Stability and Uniqueness in Signaling Games,” *Journal of Economic Theory*, 50(2), 381–413.

- CHU, C., P. LESLIE, AND A. SORENSEN (2011): “Bundle-Size Pricing as an Approximation to Mixed Bundling,” *American Economic Review*, 101(1), 263–303.
- CILIBERTO, F., AND E. TAMER (2009): “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 77(6), 1791–1828.
- CREMER, J., Y. SPIEGEL, AND C. Z. ZHENG (2009): “Symposium on Auction Theory: Auctions with Costly Information Acquisition.,” *Economic Theory*, 38(1), 41–72.
- DANIEL, K., AND D. HIRSHLEIFER (1998): “A Theory of Costly Sequential Bidding,” *University of Michigan Business School Working Paper No. 98028*.
- FISHMAN, M. (1988): “A Theory of Preemptive Takeover Bidding,” *RAND Journal of Economics*, 19(1), 88–101.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. Cambridge, Mass. and London: MIT.
- GENTRY, M., AND T. LI (2010): “Partial Identification in First-Price Auctions with Selective Entry,” *Working Paper*.
- HAILE, P. (2001): “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, 91(3), 399–427.
- HAILE, P., AND E. TAMER (2003): “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 111(1), 1–51.
- HORNER, J., AND N. SAHUGUET (2007): “Costly Signalling in Auctions,” *Review of Economic Studies*, 74, 173–206.
- KILGORE, M., R. BROWN, J. COGGINS, AND C. PFENDER (2010): “A National Review of State Timber Sale Programs and An Analysis of Factors Influencing Minnesota State Stumpage Prices,” Report to the Minnesota Forest Resources Council.
- LEVIN, D., AND J. SMITH (1994): “Equilibrium in Auctions with Entry,” *American Economic Review*, 84, 585–599.
- LI, T., AND I. PERRIGNE (2003): “Timber Sale Auctions with Random Reserve Prices,” *Review of Economics and Statistics*, 85(1), 189–200.
- LI, T., AND X. ZHENG (2009): “Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions,” *Review of Economic Studies*, 76, 1397–1429.
- (2010): “Information Acquisition and/or Bid Preparation: A Structural Analysis of Entry and Bidding in Timber Sale Auctions,” *Working Paper*.
- MAILATH, G. (1987): “Incentive Compatibility in Signaling Games with a Continuum of Types,” *Econometrica*, 55(6), 1349–1365.

- MARMER, V., A. SHNEYEROV, AND P. XU (2010): “What Model for Entry in First-Price Auctions? A Nonparametric Approach,” *Working Paper*.
- MCAFEE, R. P. (2002): “Coarse Matching,” *Econometrica*, 70(5), 2025–2034.
- MCAFEE, R. P., AND J. MCMILLAN (1988): “Search Mechanisms,” *Journal of Economic Theory*, 44(1), 99–123.
- MEAD, W., M. SCHNIEPP, AND R. WATSON (1981): “The Effectiveness of Competition and Appraisals in the Auction Markets for National Forest Timber in the Pacific Northwest,” Washington, DC: U.S. Department of Agriculture, Forest Service.
- MILGROM, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- MILGROM, P., AND J. ROBERTS (1982): “Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis,” *Econometrica*, 50(2), 443–459.
- NEEMAN, Z. (2003): “The Effectiveness of English Auctions,” *Games and Economic Behavior*, 43, 214–238.
- PAARSCH, H. (1997): “Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbia Timber Sales,” *Journal of Econometrics*, 78, 333–357.
- RAMEY, G. (1996): “D1 Signaling Equilibria with Multiple Signals and a Continuum of Types,” *Journal of Economic Theory*, 69, 508–531.
- REVLON V MCANDREWS & FORBES HOLDINGS, I. (1986): Del. Supr., 506 A.2d 173.
- ROBERTS, J., AND A. SWEETING (2011): “Competition versus Auction Design,” *Working Paper*.
- ROGERSON, W. (2003): “Simple Menus of Contracts in Cost-Based Procurement and Regulation,” *American Economic Review*, 93(3), 919–926.
- SAMUELSON, W. (1985): “Competitive Bidding with Entry Costs,” *Economic Letters*, 17, 53–57.
- SEIM, K. (2006): “An Empirical Model of Firm Entry with Endogenous Product-Type Choices,” *RAND Journal of Economics*, 37(3), 619–640.
- SUBRAMANIAN, G. (2008): “Go-Shops vs. No-Shops in Private Equity Deals: Evidence and Implications,” *The Business Lawyer*, 63(3), 729–760.
- WASSERSTEIN, B. (2000): *Big Deal: Mergers and Acquisitions in the Digital Age*. Warner Business Books.
- WILSON, R. (1987): *New Palgrave: A Dictionary of Economics* chap. Bidding. MacMillan Press.
- YE, L. (2007): “Indicative Bidding and a Theory of Two Stage Auctions,” *Games and Economic Behavior*, 58, 181–207.

A Conditions for Unique Sequential Equilibrium Under the D1 Refinement

We now verify the three conditions for our equilibrium to be the unique sequential equilibrium under the D1 refinement.

(1) $\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} > 0$. An increase in the signal threshold keeps out more second round potential entrants. The bidding behavior of those entrants who have signals above the threshold is unchanged and so an increase in S'_2 must strictly raise the incumbent's probability of winning and lower the expected price paid.

(2) $\frac{\partial \pi_v(b_1, S'_2)}{\partial b_1} / \frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2}$ is monotonic in v . Differentiating $\frac{\partial \pi_v(b_1, S'_2)}{\partial b_1} / \frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2}$ gives

$$\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial b_1 \partial v} \left(\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} \right)^{-1} - \left(\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial S'_2 \partial v} \right) \left(\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} \right)^{-2} \frac{\partial \pi_v(b_1, S'_2)}{\partial b_1} \quad (5)$$

Monotonicity requires that this expression is either always positive or always negative. We show that it is always positive by establishing (a)-(d) below.

(a) $\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial b_1 \partial v} = 0$. Consider two types of first round bidders v_H and v_L , $v_H > v_L$, with each considering increasing their bid b_1 to $b_1 + \varepsilon$, $\varepsilon > 0$. If the second bidder stays out then the change in profits for each first round type is the same, $-\varepsilon$. We now show that if the second round bidder enters the profit is still the same to each type of first round bidder.

Consider three cases. (i) $v_2 < b_1$. The first round bidder will pay $b_1 + \varepsilon$ whatever his value. (ii) $v_2 > b_1 + \varepsilon$. The final price will equal the value of the lower-valued firm and will not depend on the first round bid.³³ (iii) $b_1 \leq v_2 \leq b_1 + \varepsilon$. The first round bidder still wins, regardless of type, but now he has to pay more since before he would have won at a price of v_2 but now he wins at a price of $b_1 + \varepsilon$, yielding the same cost of $b_1 + \varepsilon - v_2$ to each type of first round bidder. Therefore, the cost of raising the deterring bid, all else constant, is independent of the first bidder's value.

(b) $\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial S'_2 \partial v} > 0$. To show that the benefit of increasing the signal entry threshold is

³³There are three cases within this case. The first is when $v_2 > v_H > v_L$. Regardless of deterring bid, both first round types would lose and so increasing the deterring bid has no effect on their profit. The second is when $v_H > v_2 > v_L$. Here the low type was going to lose regardless, and so it has no effect on his cost. Here the high type was going to win but pay v_2 no matter what and so increasing the bid has no effect on his cost. The third is when $v_H > v_L > v_2$. In either case both types were going to win but have to pay v_2 and so increasing the bid had no effect on either types' costs.

greater the higher is the first bidder's value, we can show that the benefit of excluding any second bidder type v_2 is greater, the higher is the first bidder type, regardless of v_2 . Consider the value of excluding a second round bidder whose value is v_2 for any two types of first round bidders v_H and v_L , $v_H > v_L$, both using deterring bid b_1 . If $v_2 \leq b_1$ there is no change in benefit from exclusion for either first bidder type. If $b_1 < v_2$ there are three cases. (i) $v_2 \leq v_L < v_H$. In this case the benefit of excluding the second round bidder is $v_2 - b_1$ for each first round bidder type. (ii) $v_L < v_2 \leq v_H$. In this case the benefit of exclusion is $v_L - b_1$ for the low type and $v_2 - b_1$ for the high type. Since by assumption $v_2 > v_L$, the benefit of exclusion is greater for the higher type. (iii) $v_L < v_H < v_2$. In this case the benefit of exclusion is $v_L - b_1$ for the low type and $v_H - b_1$ for the high type and so the benefit is greater for the higher first bidder type. Therefore, the benefit of excluding more second round bidders is greater the higher is the first round bidder's value.

(c) $\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} > 0$. This was shown above when we verified condition (1).

(d) $\frac{\partial \pi_v(b_1, S'_2)}{\partial b_1} < 0$. Increasing the bid is costly when it does not affect the second round potential entrant's decision. In particular, it reduces a firm's payoff when the second round firm does not enter or it enters and has a value less than b_1 . If the potential entrant enters with a value above b_1 then changing b_1 has no effect.

Combining (a)-(d), we conclude that, for all v , b_1 and S'_2 :

$$\underbrace{\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial b_1 \partial v} \left(\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} \right)^{-1}}_{=0} - \underbrace{\left(\frac{\partial^2 \pi_v(b_1, S'_2)}{\partial S'_2 \partial v} \right) \left(\frac{\partial \pi_v(b_1, S'_2)}{\partial S'_2} \right)^{-2} \frac{\partial \pi_v(b_1, S'_2)}{\partial b_1}}_{<0} > 0 \quad (6)$$

and so the monotonicity condition is satisfied.

(3) S'_2 is uniquely defined for any belief about the first potential entrant's value, and the potential entrant's response is more favorable to the incumbent when the potential entrant thinks that the incumbent's value is higher. This is true since S'_2 is a continuous function of the second period potential entrant's belief about the incumbent's value (reflecting the zero profit condition, as in equation 4, and the potential entrant's beliefs about his own value as a function of its signal) and the second potential entrant will increase S'_2 if he believes bidder 1's type is higher because his expected profits are decreasing in bidder 1's type for any signal he receives.

B Details of Estimation Method: NOT FOR PUBLICATION

In this appendix we more fully describe our estimation procedure based on Akerberg (2009)'s method of simulated maximum likelihood with importance sampling.

This method involves solving a large number of games with different parameters once, calculating the likelihoods of the observed data for each of these games, and then re-weighting these likelihoods during the estimation of the distributions for the structural parameters. This method is attractive when it is believed that the parameters of the model are heterogeneous across auctions and it would be computationally prohibitive to re-solve the model (possibly many times in order to integrate out over the heterogeneity) each time one of the parameters changes.³⁴

To apply the method, we assume that the parameters are distributed across auctions according to the specification given in Section 4.3. These specifications reflect our assumptions that σ_V , α and K are the same for mills and loggers within any particular auction, even though they may differ across auctions. The lower bound on σ_{V_a} is set slightly above zero simply to avoid computational problems that were sometimes encountered when there was almost no dispersion of values. Our estimated specifications also assume that the various parameters are distributed independently across auctions. This assumption could be relaxed, although introducing a full covariance matrix would significantly increase the number of parameters to be estimated and, when we have tried to estimate these parameters, we have not found these coefficients to be consistently significant across specifications. The set of parameters to be estimated are $\Gamma = \{\beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \omega_{\mu, \text{logger}}^2, \omega_{\mu, \text{diff}}^2, \omega_{\sigma_V}^2, \omega_{\alpha}^2, \omega_K^2\}$, and a particular draw of the parameters $\{\mu_{a, \text{logger}}, \mu_{a, \text{mill}}, \sigma_{V_a}, \alpha_a, K_a\}$ is denoted θ .

Denoting the outcome for an observed auction by y_a , the log-likelihood function for a sample of A auctions is

$$\sum_{a=1}^A \log \left(\int L_a(y_a|\theta) \phi(\theta|X_a, \Gamma) d\theta \right) \quad (7)$$

where $L_a(y_a|\theta)$ is the likelihood of the outcome y in auction a given structural parameters θ , $\phi(\theta|X_a, \Gamma)$ is the pdf of the parameter draw θ given Γ , our distributional assumptions, the unique equilibrium strategies implied by our equilibrium concept and auction characteristics including the number of potential entrants, the reserve price and observed characteristics X_a .

Unfortunately, the integral in (7) is multi-dimensional and cannot be calculated exactly.

³⁴Bajari, Hong, and Ryan (2010) use a related method to analyze entry into a complete information entry game with no selection.

We follow Akerberg by recognizing that

$$\int L_a(y_a|\theta)\phi(\theta|X_a, \Gamma)d\theta = \int L_a(y_a|\theta)\frac{\phi(\theta|X_a, \Gamma)}{g(\theta|X_a)}g(\theta|X_a)d\theta \quad (8)$$

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

$$\frac{1}{S} \sum_s L_a(y_a|\theta_s) \frac{\phi(\theta_s|X_a, \Gamma)}{g(\theta_s|X_a)} \quad (9)$$

where θ_s is one of S draws from $g(\theta|X_a)$. Critically, this means that we can calculate $L_a(y_a|\theta_s)$ for a given set of S draws that do not vary during estimation, and simply change the weights $\frac{\phi(\theta_s|X_a, \Gamma)}{g(\theta_s|X_a)}$, which only involves calculating a pdf when we change the value of Γ rather than re-solving the game.

This simulation estimator will only be accurate if a large number of θ_s draws are in the range where $\phi(\theta_s|X_a, \Gamma)$ is relatively high, and, as is well known, simulated maximum likelihood estimators are only consistent when the number of simulations grows fast enough relative to the sample size. We therefore proceed in two stages. First, we estimate Γ using $S = 2,500$ draws, where $g(\cdot)$ is a multivariate uniform distribution over a large range of parameters which includes all of the parameter values that are plausible. Second, we use these estimates $\hat{\Gamma}$ to repeat the estimation using a new importance sampling density $g(\theta|X_a) = \phi(\theta_s|X_a, \hat{\Gamma})$ with $S = 500$ per auction. Roberts and Sweeting (2011) provide Monte Carlo evidence that the estimation procedure works well even for smaller values of S .

To apply the estimator, we also need to define the likelihood function $L_a(y_a|\theta)$ based on the data we observe about the auction's outcome, which includes the number of potential entrants of each type, the winning bidder and the highest bids announced during the open outcry auction by the set of firms that indicated that they were willing to meet the reserve price. Two problems arise when interpreting these data. First, a bidder's highest announced bid in an open outcry auction may be below its value, and it is not obvious which mechanism leads to the bids that are announced (Haile and Tamer (2003)). Second, if a firm does not know its value when taking the entry decision, it may learn (after paying the entry cost) that its value is less than the reserve price and so not submit a bid.

We therefore make the following assumptions (Roberts and Sweeting (2011) present estimates based on alternative assumptions about the data generating process that deliver similar results) that are intended to be conservative interpretations of the information that is in the data: (i) the second highest observed bid (assuming one is observed above the re-

serve price) is equal to the value of the second-highest bidder; ³⁵ (ii) the winning bidder has a value greater than the second highest bid; (iii) both the winner and the second highest bidder entered and paid K_a ; (iv) other firms that indicated that they would meet the reserve price or announced bids entered and paid K_a and had values between the reserve price and the second highest bid; and, (v) all other potential entrants may have entered (paid K_a) and found out that they had values less than the reserve, or they did not enter (did not pay K_a). If a firm wins at the reserve price we assume that the winner's value is above the reserve price.

³⁵Alternative assumptions could be made. For example, we might assume that the second highest bidder has a value equal to the winning bid, or that the second highest bidder's value is some explicit function of his bid and the winning bid. In practice, 96% of second highest bids are within 1% of the high bid, so that any of these alternative assumptions give similar results. We have computed some estimates using the winning bid as the second highest value and the coefficient estimates are indeed similar.