



Selective entry and auction design[☆]

Andrew Sweeting^{a,*}, Vivek Bhattacharya^b

^a University of Maryland, United States

^b MIT, United States



ARTICLE INFO

Available online 2 April 2015

JEL classification:

D44

L10

L13

Keywords:

Auctions

Market entry

Selection

Bid preferences

ABSTRACT

This article examines how different auction designs perform when entry is endogenous and selective, by which we mean that bidders with higher values are more likely to enter. In a model where potential bidders are symmetric, we show that three alternative designs can significantly outperform the 'standard auction with simultaneous and free entry' when entry is selective. When bidders are asymmetric, we show that level of bid preference that maximizes a seller's revenues is significantly affected by the degree of selection. We also describe recent empirical and econometric work that shows that the degree of selection can be identified and estimated using standard types of auction data.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

This article examines the performance of different auction designs in a setting where bidders have independent private values but entry is endogenous and possibly *selective*. We will say that entry is selective when potential bidders with higher values are more likely to enter, as should happen when potential bidders have some information about their values prior to taking the entry decision. While it seems intuitive that entry should typically be selective, this has been ruled out by assumption in much of the theoretical and empirical auction literatures. In this article we will illustrate that allowing for selection can significantly affect the conclusions that a researcher would draw about the value of auction designs that differ from the standard auction with free entry, and that the exact degree of selection can also affect which design performs best. We also describe recent work that shows that the degree of selection is identified and can be estimated using types of data that are usually available in auction settings.

To develop our results, we consider an auction for a single unit of a good and assume that there is a well-defined set of risk-neutral potential bidders with independent private values. Throughout the article we will use the term 'player' to refer to a potential bidder, using 'bidder' to refer to a player that actually enters the auction and is able to submit a bid. The winning bidder is the one that is allocated the good at the end of

the auction. We assume that it is costly for a player to enter the auction which she must do to submit a bid. We assume that a player learns her value of the object when she incurs the entry cost, so it is natural to interpret the entry cost as including the cost of doing research or 'due diligence' on the object being sold.¹

As long as the entry cost is not too low or too high, entry will be endogenous in the sense that a player's entry decision will depend on what it expects other players to do, as well as what it believes about its own value. We model a player's belief about its value by assuming that, prior to taking the entry decision, it receives, for free, a signal that is positively correlated with its value. In equilibrium, players with signals above some threshold will enter, and the degree of correlation will, therefore, control the extent to which entry is selective. This provides us with a framework where we can examine how the degree of selection, determined by the informativeness of the signals, affects the absolute and relative performance of different mechanisms, measured either in terms of the seller's revenues or total surplus. The common assumption of no selection would involve players receiving no signals or, equivalently, signals that are completely uninformative.

We use a particular parameterization of our model to compare the performance, both in terms of revenues and efficiency, of different auction designs. Our baseline design is a 'standard auction with simultaneous and free (i.e., unrestricted) entry' (SASFE), which is the usual way that real-world auctions with endogenous entry are modeled. As

[☆] We thank Diego Moreno, several other participants at the 'Selling Mechanisms' session of the 2014 EARIE conference in Milan, Heski Bar-Isaac and two referees for comments. We are grateful to Heski Bar-Isaac for inviting Andrew Sweeting to participate in the session. Much of this research was performed in collaboration with James Roberts, and it has been financially supported by NSF grant SES-1260876. All errors are our own.

* Corresponding author.

E-mail address: sweeting@econ.umd.edu (A. Sweeting).

¹ We will assume that a player has to incur the entry cost even if she discovers that her value is less than the seller's reserve price and so does not actually submit a bid. Therefore, at least when there is a positive reserve price, it is more appropriate to interpret the entry cost as the cost of gathering information, rather than some bureaucratic cost of submitting a bid.

has been documented in the literature, one feature of this design that can be both inefficient and harmful for the seller is that entry decisions are not coordinated across players so that the realized number of bidders will be random. With symmetric players, we compare the SASFE with three designs that deal with this problem in different ways. In the ‘entry rights auction’ (ERA) of Ye (2007), the seller fixes the number of entrants in advance and conducts an initial auction for these entry ‘slots’ where players can bid based on their signals. We also consider two designs where players take entry decisions sequentially, which also allows for coordination, but also allows for the number of entrants to depend on the information that the players have and the seller does not. In one of these designs (a ‘sequential entry auction’), players decide to enter sequentially but the entrants bid simultaneously. In the other design, the ‘sequential bidding auction’ of Bulow and Klemperer (2009) (BK hereafter), players make entry decisions sequentially and entrants can submit bids when they enter in order to signal information about their values to players that are taking entry decisions later in the sequence.

When there is no selection, as is typically assumed in the literature, the SASFE generates higher expected revenues than either of the sequential designs and its revenues are quite close to those of the ERA. However, once selection is introduced into the model, both sequential designs, and especially the sequential bidding auction, generate substantially higher revenues than the SASFE and the revenue advantage of the ERA over the SASFE also increases. The identity of the mechanism that performs best depends on the exact degree of selection that is assumed. The alternative designs generate higher total surplus than the SASFE whatever is assumed about selection, but the size of their advantage over the SASFE also tends to increase when entry is more selective. In our comparisons, we draw on results developed in Roberts and Sweeting (2013), who compare the SASFE and the sequential bidding auction, and Bhattacharya et al. (2014), who compare a SASFE and an ERA in a procurement setting. The new results in the present article come from using a single set of parameters, so that the sequential bidding auction and the ERA can also be compared; adding the sequential entry auction to the consideration set; considering how absolute and relative performance changes when we move from no selection to partial selection (the earlier papers only consider different degrees of partially selective entry); and, examining in more detail why the alternative mechanisms are more efficient and generate higher revenue.

We also examine how the degree of selection in the entry process can affect the performance of bid preference programs, that are widely used by government agencies when selling assets or procuring services, in a model where bidders are asymmetric.² These programs are partly motivated by wanting to increase the probability that bidders of a particular type will win, but also, following the logic of optimal auctions (Myerson, 1981), by a desire to raise auction revenues by increasing the competition that strong bidders face. We show that while very large bid preferences maximize revenues when there is little selection, much smaller preferences are optimal when the degree of selection is high. These results are also new, and an additional contribution is that we use our analysis to illustrate how changing the degree of selection changes the level of entry costs required to rationalize a given amount of entry by weak bidders.

Our paper contributes to the enormous theoretical literature on auction design, summarized in the surveys of Klemperer (2004), Krishna (2002) and Milgrom (2004). When the seller has a single unit, and there is a fixed number of risk-neutral and symmetric bidders with independent private values, it is well-known that the optimal mechanism

is a standard auction with a reserve price or entry fee.³ Much of the theoretical auction literature has been concerned with exploring which mechanisms perform best when these assumptions are relaxed. We will focus on relaxing the assumption that the number of bidders is exogenous, and explore how particular assumptions about the way the entry process works affect the absolute and relative performance of different mechanisms. Milgrom (2004) uses endogenous entry as his leading illustration of why auctions need to be analyzed in their correct context, arguing that even auctions that are carefully designed can fail when too few bidders decide to participate (p. 209).⁴

We follow the existing literature in modeling the way that standard auctions work as a two-stage game, where, in the first stage, players simultaneously decide whether to enter, incurring a common entry cost, and, in the second stage, the entrants simultaneously submit bids. This is what we will label an SASFE. Entry decisions into an SASFE will be non-trivially endogenous when the entry cost is ‘moderate’ (Milgrom, p. 217) in the sense that it is low enough that, in equilibrium, some players may want to enter, while being high enough that some may not.

The cleanest set of theoretical results come from models in which players have no private information about their values until they have entered, so that entry is not selective. Assuming that players are symmetric, that the common entry cost is moderate and that the entry game is followed by a standard first price or second price auction (revenue equivalence holds in this context), Levin and Smith (1994) show that (i) the symmetric equilibrium involves players mixing over whether to enter, and making zero expected profits; (ii) the seller’s optimal reserve price is equal to its value of keeping hold of the object, with revenue-maximization requiring no reserve price and no entry fees (see also McAfee and McMillan, 1987); (iii) an increase in the number of potential entrants will reduce expected revenues; and, (iv) when the reserve price is equal to the seller’s value, equilibrium entry strategies are optimal in the sense that a social planner who also had to choose a symmetric entry rule would choose the same entry probability that the players themselves choose in equilibrium. In what follows, we will refer to the assumption that entry is not selective as “NS”.⁵ Of course, property (iv) does not imply that the mechanism is necessarily optimal when compared to mechanisms where the seller changes the entry process in some way, such as fixing the number of players that can enter or organizing players to move sequentially.

Assuming NS, BK compare outcomes in a SASFE with those in an alternative procedure where players take entry decisions and bid sequentially, which they argue is a stylized version of how corporations are often sold. They show that the alternative procedure raises total surplus but will almost always generate lower revenues for the seller, because of the ability of early movers to deter entry. We will show that their sequential bidding procedure can actually increase revenues quite significantly as soon as any degree of selection is introduced into the model.

A more limited literature has considered endogenous entry with selection. Samuelson (1985) and Menezes and Monteiro (2000) assume that players know their values when deciding whether to enter. This is the most extreme form of selection that we will consider, and we will call this the fully selective, “FS” assumption. A feature of this model is that bidders with high values tend to make positive profits in equilibrium. In the SASFE under FS, revenues may increase or decrease when additional players are added, and the seller-optimal reserve may be greater than the seller’s value of holding onto the object

³ Bulow and Klemperer (1996) show that under these assumptions, adding an additional bidder in a standard auction will increase the seller’s revenue by more than using the optimal design with a fixed number of bidders (which involves setting a reserve price). As our results illustrate, this conclusion does not necessarily hold when entry is endogenous and one considers the effects of adding a *potential* bidder.

⁴ Milgrom’s second illustration concerns asymmetries between bidders, which we also consider.

⁵ The “not selective” assumption is sometimes called the “LS” assumption after Levin and Smith. Similarly, the “fully selective” assumption that we introduce as the opposite polar case below is often referred to as the “S” assumption following Samuelson (1985).

² Roberts and Sweeting (2013) allow for players to be asymmetric in the context of a second-price auction, while Bhattacharya et al. (2014) consider a low-bid auction with symmetric players. In the current article, we show that it is feasible to solve first-price auctions with asymmetric bidders and selective and endogenous entry. This framework is appropriate because bid preference programs are usually applied in the context of first-price or low-bid auctions.

(Menezes and Monteiro, 2000). Characterization of optimal policies tends to be specific to the value distributions considered, which is one reason why the NS assumption has been the focus of most analysis. Hubbard and Paarsch (2009) use a computational approach to consider the effects of bid preferences to some subset of symmetric players in first-price auctions under the FS assumption.⁶

Ye (2007), Bhattacharya et al. (2014), Roberts and Sweeting (2013), Marmor et al. (2013) and Gentry and Li (2014) consider models where players face a common entry cost but have noisy signals about their values before they enter, as we assume in this article.⁷ Moreno and Wooders (2011), Cremer et al. (2009) and Lu and Ye (2013) consider a variant of the NS model where players have heterogeneous entry costs. We consider below whether this model has similar implications to one where entry is partially selective.

Most empirical work has also tended to make the NS assumption, partly because this assumption implies that the distribution of players' values will be the same as the distribution of bidders' values which is what can be inferred from the data. Athey et al. (2013) and Krasnokutskaya and Seim (2011) both consider the effects of bid preferences when entry is endogenous but bidders do not know their values when deciding whether to enter. In Section 6 we show that smaller bid preferences may be optimal for the seller when entry is selective. In Section 7, we describe recent results showing that selective entry models are identified, as well as noting recent empirical work, including Roberts and Sweeting (2013), Roberts and Sweeting (2015) (timber) and Bhattacharya et al. (2014) (highway procurement), that show that these models, including the degree of selection, can also be estimated in practice.

As models with endogenous and selective entry are not analytically tractable, our comparisons are computational. For ease of exposition, we focus on a single set of parameters for most of the analysis. While our results hold for a wide variety of parameters and value distributions that we have considered in our research, the reader should be clear that we are only claiming that selective entry can matter for the relative performance of different mechanisms, not that it must always do so. We do not compare the various mechanisms that we consider with optimal mechanisms, partly because when entry is partially selective, the generally optimal mechanism is unknown although, in very recent work, Lu and Ye (2014) have characterized the optimal design of a two-stage auction.⁸

The article is structured as follows. Section 2 introduces the basic model, assuming that players are symmetric, and describes equilibrium strategies and the effects of selection when a standard auction with free entry is used. Section 3 describes the alternative mechanisms considered with symmetric bidders and their associated equilibrium strategies. Sections 4 and 5 contain the comparisons of expected total surplus (efficiency) and revenues. Section 6 considers bid preferences in a setting where players are asymmetric. Section 7 briefly describes recent work that shows that the degree of selection is non-parametrically identified and can be estimated using real world data. Section 8 concludes.

2. Model

In this section we outline the basic model that we will use to compare different auction designs and illustrate some of the properties of

the equilibrium outcomes in the SASFE. For now, we will assume that players are symmetric, leaving all discussion of the asymmetric case to Section 6.

2.1. Informational assumptions

We assume that there are N players interested in a single unit of an asset. These players have independent private values, which are i.i.d. draws from a distribution $F^V(v)$ (pdf $f^V(v)$), which is continuous on an interval $[0, \bar{V}]$. N and $F^V(v)$ are commonly known by all players and the auction designer. To be able to submit a bid for the asset, a player must incur an entry cost K . A player that incurs this entry cost (i.e., a bidder) is assumed to find out her value for sure, so a natural interpretation is that K contains a 'due diligence' cost associated with evaluating the asset although it could also include other costs of participation, such as securing the bonds that are often required in procurement auctions. In what follows we will assume that K is fixed, and is not a parameter chosen by the auction designer. Prior to incurring the entry cost, each player receives a private information, noisy signal of her value. Specifically we will consider the case where $s_i = v_i z_i$, $z_i = e^{\varepsilon_i}$, $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ and the ε s are i.i.d. across bidders, although the functional form is not important.⁹ A player is therefore able to condition her entry decision on her own signal. We assume that the seller has no value to retaining the object, so that its objective is revenue-maximization.¹⁰ The FS model corresponds to the case where $\sigma_\varepsilon^2 = 0$, so the signal is perfectly informative of the player's value. As $\sigma_\varepsilon^2 \rightarrow \infty$, we approximate the informational assumptions of the NS model, where there are no signals, although, as we shall show, this does not necessarily mean that strategies and outcomes under NS are always similar to those in a model with very uninformative signals.

2.2. Standard Auction with Simultaneous and Free Entry (SASFE)

To illustrate these assumptions, it is useful to describe equilibrium strategies in our baseline mechanism, which is the standard model used in the literature to describe most real-world auctions (*inter alia* (Levin and Smith, 1994; Athey et al., 2011; Athey et al., 2013; Krasnokutskaya and Seim, 2011; Li and Zheng, 2009; Bhattacharya et al., 2014)). In particular, there is a two-stage game where, in the first stage, players simultaneously and non-cooperatively decide whether to enter, and in the second stage, the entrants compete in a simultaneous second-price or first-price auction. As revenue equivalence holds when bidders are symmetric, we will, for tractability, formulate the second stage procedure as a second-price auction.¹¹ We allow for the possibility that the seller sets a reserve price r that is known to all players at the beginning of the game. Recall that a standard auction with a reserve price is the seller's optimal mechanism when entry is fixed.

In the second stage, all entrants are assumed to bid their values, so that the good will be allocated to the bidder with the highest value at a price equal to the second highest bidder value. In the first stage, a player will enter if her private signal exceeds a threshold determined by a

⁶ We extend their computational framework below to consider the effect of bid preferences when players are asymmetric and entry is imperfectly selective.

⁷ With symmetric players, property (iv) from the NS model continues to hold.

⁸ Cremer et al. (2009) characterize the form of the optimal mechanism under the NS assumption with possibly heterogeneous entry costs. They show that it could be implemented by a sequence of auctions with appropriate reserve prices, and possible entry subsidies and entry fees. Lu and Ye (2013) design the optimal two-stage auction for this case. Lu and Ye (2014) show the optimal form of a two-stage auction where players compete to enter an auction for the object and entry may be selective. The entry rights auction that we consider here differs from their optimal form in that the number of entrants to the second-stage is fixed in advance rather than being a function of first-stage bids, and that our second-stage does not involve those firms that are selected being handicapped based on their first stage bids.

⁹ More generally, and following the exposition in Gentry and Li (2014), one can think of each player receiving a signal that has a uniform marginal density on $[0, 1]$, where the conditional distribution of v_i is stochastically ordered in s_i . The assumption of independence across bidders is more important, as a player can only infer something about another player's value from that player's entry decision, rather than something about her own value.

¹⁰ This assumption could potentially affect the results by favoring mechanisms that increase the probability that the object is sold. However, the probability that the object is not sold in an SASFE is quite small, and ignoring these cases does not change the comparison between the different auction mechanisms.

¹¹ In this setting, revenue equivalence holds whether or not bidders are informed of the number of other bidders that entered into a first-price auction. See Section 6.2 of Milgrom (2004) for an extended discussion.

zero profit condition. Focusing, as usual, on the symmetric equilibrium, the equilibrium threshold S^* will solve

$$\int_r^{\bar{V}} \left[\int_r^v (v-x)h(x|S^*) dx \right] g(v|S^*) dv - K = 0$$

where

$$g(v|s_i) = \frac{f^V(v) \frac{1}{v\sigma_\varepsilon} \phi\left(\frac{\ln\left(\frac{s_i}{v}\right)}{\sigma_\varepsilon}\right)}{\int_0^{\bar{V}} f_r^V(x) \frac{1}{x\sigma_\varepsilon} \phi\left(\frac{\ln\left(\frac{s_i}{x}\right)}{\sigma_\varepsilon}\right) dx}$$

is the pdf of values conditional on signal s_i , $\phi(\cdot)$ denotes the standard normal pdf and $h(x|S^*)$ is the conditional pdf of the highest value of other entering bidders (or r if no other firm enters with a value above the reserve price) when they also use entry threshold S^* . It is straightforward to show that there will be a unique symmetric equilibrium, and we will assume that this equilibrium is played, although asymmetric equilibria may exist.

The variance of the signal noise has a significant effect on auction efficiency and the distribution of payoffs between buyers and sellers, and this has important implications for auction design. To illustrate this, we introduce the parameters that we will use throughout our analysis of sales to symmetric players. We assume that $\ln(V_i) \sim N(\mu_V, \sigma_V^2)$, with V truncated to $[0,200]$, $\mu_V = 4.5$, $\sigma_V = 0.2$, and that $N = 5$ and $K = 5$ (which is moderate - i.e., in equilibrium, some players may want to enter while others may not - for all assumptions on the degree of selection). Unless otherwise stated, we assume that the reserve price $r = 0$. The mean value of a player is 91.8, with standard deviation 18.5.¹² The truncation point, $\bar{V} = 200$, is sufficiently high that the probability that a value from the untruncated log-normal would lie above this value is small, and our calculations are essentially unaffected if we consider even higher values for this parameter.

A useful measure of the precision of the signal is $\alpha \equiv \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_V^2}$, which must lie on $[0,1]$.¹³ Fig. 1 shows the density of a player's posterior belief about her value when she receives a signal equal to the 75th percentile of the marginal distribution of signals for four different values of α . For $\alpha = 0.9$ or 0.99 , the conditional distributions are similar to the marginal distribution of values (also shown), which would, of course, be a player's belief about its value in the NS model. In spite of this, we will show that auction performance can change quite significantly when one moves from the NS model to a selective entry model with a fairly high value of α , such as 0.9. It is in this sense that we will claim that "small deviations" from the NS assumption can matter, although one might also view our finding that outcomes can differ substantially as implying that the differences between a NS model, where players have no private information when they decide whether to enter, and one where they have very noisy signals, are not really as small as Fig. 1 might suggest.¹⁴ Lower values of α are associated with more precise signals,

¹² These parameter choices are somewhat arbitrary, but it seems plausible that the distribution of values will generally be single-peaked and that there should be reasonable heterogeneity in values for assets that sellers decide to sell by auction.

¹³ If the value distribution were not truncated above, the posterior distribution for a player's value having received a signal s_i would be a log-normal distribution with location parameter $\alpha\mu + (1 - \alpha)\ln(s_i)$ and squared scale parameter $\alpha\sigma_V^2$, so one can think of $1 - \alpha$ as, approximately, the weight that the player should place on her signal when updating her beliefs.

¹⁴ This might matter because, for example, an applied researcher might use the fact that due diligence work takes place as a reason to prefer the NS model over the FS model. However, due diligence would still be required even if, as seems very plausible in almost all settings in practice, players have some information about their private values before they begin the due diligence process. Our results imply that conclusions derived under NS may not hold if some selection was allowed.




Fig. 1. Marginal distribution of values and the posterior distribution of values conditional on receiving the 75th percentile signal.

so that for a given entry threshold, a potential bidder with a signal above that threshold is more likely to have a high value. We will, therefore, say that a lower α is associated with more selective entry.

Table 1 shows how a set of expected outcomes, specifically, seller revenues, the value of the winning bidder, the number of entrants, total surplus (measured as the value of the winning bidder less total entry costs) and bidder profits, vary as a function of α . The bottom row shows the results when players receive no signals and, in equilibrium, mix over whether to enter or not, as in the NS model.

In the SASFE, outcomes under NS are quite similar to those when $\alpha = 0.99$, but they change quite quickly as α falls, so that roughly half of the decrease in revenues and the increase in bidder profits which happens when one moves from the NS case to the FS case occurs

Table 1
Expected Outcomes in a Standard Auction with Simultaneous and Free Entry.

α	Value of Winner	Number of Entrants	Total Surplus	Seller Revenues	Bidder Profits
0 (FS)	112.29	2.57	99.44	85.57	13.88
0.01	112.29	2.57	99.44	85.57	13.87
0.1	112.25	2.60	99.27	85.62	13.65
0.2	112.22	2.63	99.06	85.77	13.30
0.3	112.14	2.67	98.78	85.96	12.81
0.4	111.98	2.72	98.38	86.16	12.22
0.5	111.78	2.77	97.92	86.45	11.47
0.6	111.55	2.83	97.37	86.78	10.59
0.7	111.25	2.91	96.70	87.22	9.47
0.8	110.86	3.00	95.82	87.77	8.05
0.9	110.32	3.14	94.60	88.58	6.01
0.95	109.96	3.25	93.69	89.24	4.44
0.99	109.58	3.42	92.45	90.30	2.14
NS	109.43	3.60	91.42	91.42	0

Notes: When there is no entry, the value of the winner, revenues and surplus are zero. Bidder profits are the sum of player profits. Expected outcomes based on 500,000 simulations.

when one moves from NS to $\alpha = 0.9$. This finding will be common to many of our comparisons, even though players' posteriors under NS are not so different to those when $\alpha = 0.9$, as illustrated in Fig. 1.

Several other patterns are also interesting. First, when entry is more selective, there is less entry and, therefore, lower spending on entry costs.¹⁵ However, because potential entrants are more informed about their values, the expected value of the winning bidder actually tends to increase even though there is less entry and the distribution of the highest value in the population of players is unchanged. Putting these two features together means that expected surplus increases quickly as selection increases.

Second, selection has a quite dramatic effect on the distribution of surplus. Under the NS model, the fact that players mix over entry in equilibrium implies that entrants' expected profits are equal to zero. Therefore, in expectation, the seller captures all of the surplus. On the other hand, with selection only the marginal entrant has expected profits of zero, while inframarginal entrants, with more optimistic posterior beliefs about their values, expect positive profits. Fig. 2 illustrates the difference between the distribution of values for the marginal (i.e., the pdf of values conditional on receiving the signal $s = S^*$) and average inframarginal entrant (the pdf of values conditional on receiving a signal $s \geq S^*$) for two different values of α . When selection is weak ($\alpha = 0.9$) these distributions are much more similar than when entry is selective.

This is an appropriate place to comment briefly on the inefficiencies that exist in the entry process of the SASFE model, and to clarify some of the discussion of the inefficiencies that exists in the literature. As explained by Milgrom (2004), in the context of the NS model, and Gentry and Li (2012), in the context of a selective entry model, entry strategies are efficient in the sense that a social planner that was constrained to choose, *ex-ante*, an identical entry probability or threshold rule for all players would choose the same strategies that the players themselves would choose in the symmetric equilibrium.¹⁶ However, this does not mean that entry decisions are efficient in a more general sense. For example, when all potential entrants make simultaneous decisions using the same thresholds, the realized number of entrants is random, and when surplus is a concave function of the number of entrants (as can be easily shown under NS), surplus can be increased by fixing the number of entrants.

One way to fix the number of entrants, which makes particular sense when players have some private information about their values, is to hold an auction for a limited number of slots to compete in an auction for the object. This is the "entry rights auction" that we consider below. An alternative way to try to address some of the inefficiency that arises from randomness without fixing the number of entrants is to make players take their entry decisions sequentially rather than simultaneously. We consider two sequential procedures below: in one of them players enter sequentially but bid simultaneously; and, in the other one, there is also an element of sequential bidding. This opens up the possibility that subsequent entry decisions may be conditioned, to some extent, on the values, as well as the entry decisions, of earlier-movers. This may be more efficient, but increases the possibility that earlier-movers will be able to deter later entry, hurting the seller.

¹⁵ In general, we would expect less entry with selection, when K is held fixed, for two reasons. First, holding the probability that other players enter fixed, those entrants will tend to have higher values when entry is more selective decreasing the expected payoff that a player with a given value has from entering. Second, a player's surplus in an second-price auction is a convex function of her own value. Without selection, the expected surplus is calculated using the unconditional distribution of these values. With selection, it is calculated using a distribution that is conditional on the signal received, and, because this distribution should be more concentrated than the unconditional distribution, the expected payoff from entering will, all else equal, tend to be lower.

¹⁶ In a second price auction for a single unit a Mankiw and Whinston (1986) style 'excess entry' result does not hold because an entrant only takes market share from other firms when it is socially efficient to do so.

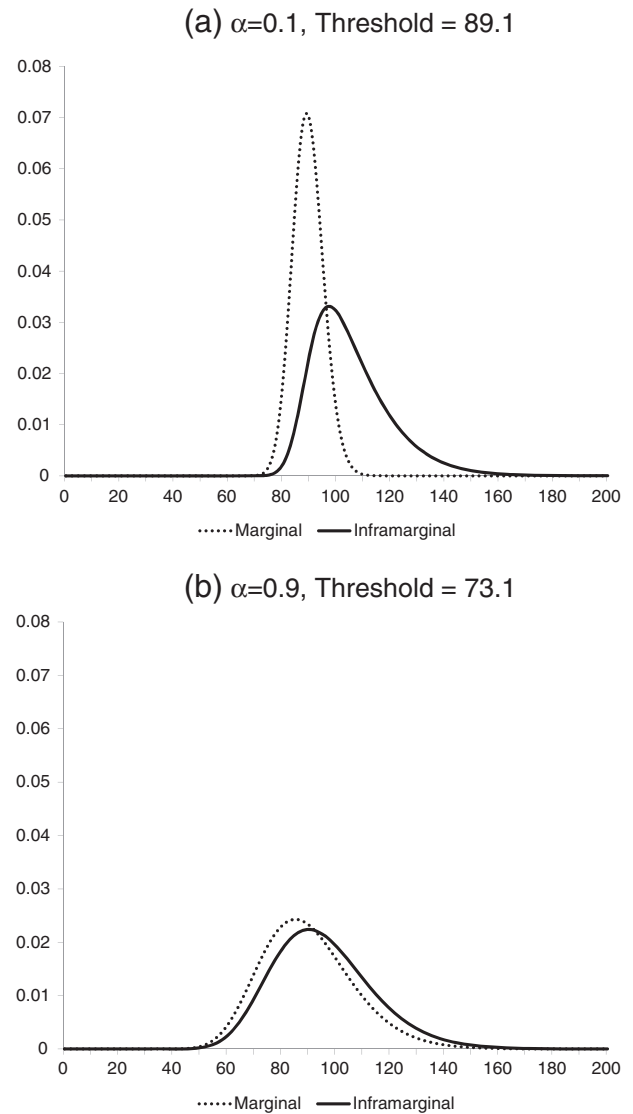


Fig. 2. Distribution of values for the marginal and the average inframarginal entrant for different degrees of selection.

An alternative model of the standard auction, which has been considered in Moreno and Wooders (2011) and which has been used empirically by Krasnokutskaya and Seim (2011) and Li and Zheng (2009), is to assume that entry is simultaneous and non-selective, but that players have heterogeneous, privately observed entry costs. The equilibrium entry strategies in this case are also threshold rules, in entry costs instead of signals, and an obvious similarity to the selective entry model is that players with low entry costs will tend to make positive profits in equilibrium. An important difference is that while a lower entry cost makes entry more attractive, it does not affect the payoffs of other players conditional on entry occurring, whereas when a player has a high signal in a selective entry model this does affect other players' expected payoffs from entering because it is the bidder with the highest value that will be allocated the object. However, this leaves open the question of whether the models have different implications for auction design.

In order to shed some light on this question, Table 2 shows the same outcomes as Table 1, as a function of σ_K , when entry costs are assumed to be distributed normally with mean 5 (the value of K assumed in our examples) and standard deviation σ_K , and we assume NS, so that players receive no signals prior to taking their entry decisions. The first row corresponds to the model in the final row of Table 1.

Table 2
Expected Outcomes in an NS Model with Heterogeneous K .

σ_K	Value of Winner	Number of Entrants	Total Surplus	Seller Revenues	Bidder Profits
0, common K (NS)	109.43	3.60	91.42	91.42	0
0.01	109.42	3.60	91.45	91.38	0.07
0.05	109.38	3.59	91.51	91.30	0.22
0.1	109.34	3.58	91.60	91.20	0.40
0.2	109.24	3.56	91.78	91.00	0.77
0.3	109.16	3.55	91.94	90.81	1.13
1	108.99	3.52	92.28	90.45	1.83
2	108.06	3.35	94.94	88.36	6.58
5	106.91	3.16	100.49	85.68	14.81

Notes: Mean K is 5. σ_K is the standard deviation.

Expected outcomes based on 500,000 simulations.

Increasing the heterogeneity in entry costs raises total surplus and bidder profits, while reducing entry and seller revenues. In terms of direction, these are the same changes that come from increasing the degree of selection. However, the causes of the changes in total surplus are somewhat different between the two models. In the selective entry model, the value of the winner tends to rise and the amount of entry tends to fall with more selection, and both of these forces raise total surplus. In contrast, with heterogeneous K , the value of the winner and the expected amount of entry must move in the same direction, and both fall when there is more heterogeneity. The increase in total surplus is driven by the fact that some entrants actually have lower entry costs, whereas, in the selective entry model neither the distribution of valuations nor the level of entry costs change when selection is introduced.¹⁷ Maybe the most striking difference, however, is that when we allow for limited heterogeneity in entry costs, which seems plausible in reality,¹⁸ the expected outcome measures only change by small amounts (for example, going from $\sigma_K = 0$ to $\sigma_K = 0.3$ lowers a seller's revenue by less than 1%), whereas seemingly small deviations from the informational assumptions of NS (e.g., moving from NS to $\alpha = 0.9$) have already been shown to change outcomes more dramatically. To generate revenues and surplus similar to the FS model, one needs to push σ_K as high as 5, by which point, given our parametric assumptions, 16% of players have negative entry costs, in which case their entry *increases* surplus even when they do not win the object.¹⁹

3. Alternative mechanisms

With symmetric bidders, we focus our discussion of mechanisms other than the SASFE on three relatively simple alternatives, two of which have been considered in the existing literature. Of course, these alternatives do not exhaust the spectrum of possible alternatives, but, to the extent that we find that alternatives outperform the SASFE, our conclusion would only be strengthened if we found other mechanisms that could do even better.

¹⁷ For example, when $\sigma_K = 2$, efficiency would be lower than in the common entry cost case if all of entrants had to incur the mean entry cost.

¹⁸ Note that in models that allow heterogeneous entry costs, it is assumed that this heterogeneity is not correlated with players' valuations, so the heterogeneity must be interpreted as being due to differences in the technology of evaluating the object or in the bureaucratic costs of submitting bids, rather than being due to the type of expertise that might be associated with having a high value. It seems unlikely that differences in costs of research technologies that players would actually choose to use would be large.

¹⁹ An open, but interesting, question concerns how parameters will be biased if a researcher estimates a model that allows for heterogeneity in entry costs, when the true model has no heterogeneity but does have selection. The existing empirical literature on market entry, including Krasnokutskaya and Seim (2011) in an auction context, has typically estimated entry costs to be quite heterogeneous, and one might conjecture that large estimated variances may actually reflect the presence of selection.

3.1. Entry rights auction (Ye, 2007; Bhattacharya et al., 2014)

Ye (2007) characterizes equilibrium strategies in a game where the seller first announces the number of bidders it will allow to compete in an auction for the object and then auctions off these rights to all players in a first-stage auction, before conducting the auction itself.²⁰ This 'entry rights auction' (ERA) procedure addresses the problem that the number of entrants into an SASFE is random, by having the seller control the number of entrants, while guaranteeing that it is the players with the highest signals that enter. It also allows the seller to extract some of surplus that the restricted set of entrants will get from the second-stage auction. However, it does not allow the number of entrants to be a function of the private information that players have, which can be a disadvantage relative to other mechanisms when entry is partially selective.

Following Ye, we assume that an ERA works in the following way. First, the seller commits to select the n highest first-stage bidders to compete in an auction for the object and to give each of the selected bidders a subsidy equal to K . Then, all potential bidders, having seen their signals, submit non-negative bids in a first-stage auction that uses an all-pay format. The n selected bidders then receive their subsidy, incur their entry cost K , find out their values, and compete in a second-stage auction for the object, which we will assume has a second-price format. There are no reserves in either auction.

Some of the details of the mechanism deserve further comment. First, following Ye's Proposition 5, the all-pay format and the subsidy, while relatively rarely observed in practice, are used to guarantee that first-stage bids are strictly increasing functions of signals,²¹ so that the selected second-stage bidders will be those that are most likely to have the highest values. Second, any subsidy that guarantees this type of efficiency will generate the same (net) revenue for the seller (Proposition 6). Third, while more standard first-stage formats in which only the selected bidders pay might not have strictly increasing first-stage bid functions, for parameters where these functions are monotonic these formats should generate the same revenue and efficiency outcomes as the all-pay format (this also follows from Proposition 6).²² For this reason, focusing on the all-pay format as a modeling device is reasonable.

In our view, the most important caveat associated with the mechanism is that it is assumed that participation in the first-stage auction is costless, whereas participation in the second-stage auction requires K to be incurred. This may be unreasonable from the perspective that there may be some bureaucratic costs associated with submitting any type of bid, even if no due diligence is done, but it might also be difficult to generate interest in the auction if potential buyers are unable to conduct some examination of the asset before they submit binding first-stage bids.²³ In practice, two factors may tend to lessen the force of this critique in some real-world settings. First, if a seller has to make frequent use of auctions it will have an incentive to develop a reputation for not holding auctions for the right to try to buy worthless objects. Second, when $n \geq 2$ the revenues from the first-stage bids are usually much

²⁰ Ye also compares this type of auction, where first-stage bids are binding and result in payments, with 'indicative bidding' schemes where bidders are only asked for indications of what they will bid in the first-stage. He shows that indicative bidding schemes generically do not have equilibria that result in efficient entry. Quint and Hendricks (2013) show that indicative schemes may have weakly monotonic first-stage equilibrium bidding strategies when the bidding space is discrete and either entry costs are large or the number of bidders is large.

²¹ For example, as Ye notes, the FS model presents a particular problem for other auction formats. In those formats a pure strategy first-stage bid would be determined only by the value that a player has to being the marginal entrant into the second stage, but when $n \geq 2$, the marginal entrant will certainly lose in the second stage under the FS assumption, so that first-stage bids will be zero for all signals.

²² Monotonicity in first-stage uniform or discriminatory formats requires that a bidder's expected payoff from being selected conditional on being the *marginal* selected entrant is increasing in her signal. For many parameters, including those used in our examples, this fails, but often only for quite high signals that occur with relatively low probability. Therefore, it is at least plausible that these standard formats might work quite well in practice, but this is obviously an open question.

²³ Note that this criticism could also apply to some types of entry fee.

smaller than the amounts that the winner will pay in the second-stage, especially when entry is very selective.²⁴ Therefore, it is plausible that firms may be willing to submit first-stage bids on the basis of much more limited due diligence that they would need to do if selected for the second-stage auction.

In equilibrium, the selected bidders will submit bids equal to their values in the second-stage, which resembles a standard independent private values second-price auction with a fixed number of bidders. To solve for the first-stage equilibrium bid function, $\gamma^*(\cdot)$, note that the expected second-stage profit of a selected entrant with value v , when the $(n + 1)^{st}$ highest signal is \bar{s} , is

$$\Pi(v; \bar{s}) = \int_0^v (v-x)h_{V|S>\bar{s}}(x)dx > 0$$

where the inequality follows from the fact that the entry cost is fully subsidized for the selected bidders and $h_{V|S>\bar{s}}$ is the pdf of the highest value of other bidders given \bar{s} . Then $\gamma^*(s)$ should solve

$$\gamma^*(s) \equiv \arg \max_g \int_0^{\gamma^{*-1}(g)} \int_0^{\bar{v}} \Pi(v; \bar{s}) dF_{V|S=s}(v) dF_{S^{(n-1)}}(\bar{s}) - g,$$

where $F_{V|S=s}(\cdot)$ is the conditional distribution of the value given a signal of s and $F_{S^{(n-1)}}(\cdot)$ is the distribution of the n^{th} highest signal of the remaining $N - 1$ bidders, with pdf $f_{S^{(n-1)}}(\cdot)$. Solving this maximization problem and imposing that, in equilibrium, $g = \gamma^*(s)$ gives the differential equation

$$\gamma^{*'}(s) = \int_0^{\bar{v}} \Pi(v; s) f_{S^{(n-1)}}(s) dF_{V|S=s}(v),$$

with boundary condition $\gamma^*(0) = 0$. As $\Pi(v; \bar{s}) > 0$, the first-stage bid function is monotonically increasing in s .

The value of n , the number of bidders selected for the second stage, can play an important role in the ERA, and is chosen by the seller. If $n = N$, which is never optimal for the level of entry costs that we consider, first-stage bids will be zero and the seller effectively subsidizes the entry costs of all potential bidders in the second-stage. If $n = 1$, then the procurer selects a winner based only on signals.²⁵ In the calculations in the next section we assume that $n \geq 2$, and find that $n = 2$ is optimal for the seller, in terms of revenues, unless α is very close to 1.

While this model is appropriate for selective entry, in the NS case, where bidders have no signals, there is no pure strategy equilibrium in the all-pay auction. Instead, for this case only, we will assume that the seller uses a uniform price first-stage auction, where only the selected bidders pay and the price is equal to the $(n + 1)^{st}$ highest bid. In this case, each player will bid its value to being selected for the second stage in the symmetric pure strategy equilibrium.²⁶ As all of these bids will be the same, entrants will, in effect, be chosen randomly.

3.2. Sequential auctions

We consider two sequential procedures. In one of them, the sequential entry auction, players choose to enter sequentially, in a random order chosen by the seller, but entrants submit simultaneous bids once all entry decisions have been made. In the other, players may place bids as soon as they enter, and we will call this version the sequential bidding auction.

3.2.1. Sequential entry auction

In this procedure, the seller approaches the players in a random order and asks them whether they wish to enter. Players later in the sequence observe earlier choices and we assume that players are committed to the choices that they make, so that a player that chooses to enter will incur the entry cost K , and find out its value, and one that chooses not to enter cannot subsequently change this decision. Once all players have taken entry decisions, firms submit bids in a second-price auction.²⁷

Under NS, the entry game has a simple equilibrium where players will enter until an additional entrant would expect negative profits from entry given the entry that has already taken place. For our parameters, for example, exactly three players enter in equilibrium. In a model with signals, equilibrium entry rules will be thresholds, and a player's threshold will depend on the decisions of earlier entrants, the thresholds that those players used and the thresholds that later players will be using. For example, if the first player has an entry threshold of 100, then a later player who observes that the first player entered can infer that the first player's signal was at least 100, and this later player's entry decision will also be affected by the thresholds that it believes players that get to move even later will use if it enters. Therefore, rather than solving for a single equilibrium threshold it is necessary to simultaneously solve for an entire vector of thresholds, based on different zero profit conditions, as a function of all possible histories of the game.²⁸ One issue that arises in these auctions is that, because later movers must form beliefs about the values of earlier players, it is possible that the entry game may have multiple equilibria even though entry is sequential. The results reported below for $N = 5$ are based on equilibria that were found from multiple starting points.²⁹

3.2.2. Sequential bidding auction (Bulow and Klemperer, 2009; Roberts and Sweeting, 2013)

In an important paper, BK seek to understand the relative performance of a standard free-entry auction and a sequential procedure where bidders are allowed to both enter and place bids in a sequential order. They motivate the latter as being a stylized model of the way that many corporations are sold. For example, the board of the target corporation may negotiate an outline deal with one potential buyer, before other potential buyers are given an opportunity to submit competing bids. BK compare the mechanisms partly to understand whether the use of sequential procedures is in the interests of target or acquiring shareholders.³⁰ Under the assumptions of the NS model, they show that the standard auction will almost always generate higher expected revenues for the seller, even though the sequential procedure, which allows later players to condition their entry decisions on how earlier-movers have both entered and bid, will tend to be more efficient. This reflects

²⁷ One might argue that if players have some sense of the entry decisions that other players are making, sequential entry might be an appropriate way to model how standard auctions actually work, although this is not the usual modeling approach. An exception is BK who use sequential entry as their baseline model of entry into the standard auction, although they note that their conclusions do not change if entry is simultaneous.

²⁸ For example, with two players there are three equilibrium entry thresholds (one for the first player, and two for the second player); seven thresholds for three players; fifteen thresholds for four players; thirty-one for five players; and so on. The computational burden therefore increases geometrically in the number of players.

²⁹ For $N = 6$ we did find multiple equilibria when entry is quite selective. These equilibria differed in some of their revenue and efficiency properties. For example, many of the equilibria involve some later players in the sequence entering almost certainly with other players almost certainly staying out, so that outcomes look more like those that would occur if a random group of bidders were selected to participate in the auction independent of their signals. In these equilibria, expected surplus tends to be lower (in some examples by as much as 3%), with smaller decreases in expected revenues. We do not report sequential entry auction results for $N = 6$, although it would be interesting to explore the number of equilibria and the implications of the various equilibria for auction design, and whether a seller might be able to help select particularly favorable equilibria using a more exotic design, in a more systematic way. As pointed out by a referee, one might view the existence of inefficient equilibria as one practical argument against using the sequential entry auction.

³⁰ The board of the target has a legal duty to act in the interests of its shareholders. Denton (2008) questions the legality of "go-shop" procedures, which have a sequential element, as a way of selling corporations, based on BK's results.

²⁴ In our example, expected total first stage revenues are highest under NS (16.98) but they are less than 5 for values of $\alpha < 0.3$.

²⁵ In the case of the FS model, $n = 1$ would be socially efficient because only the highest value entrant would incur the entry cost, and, if the seller could set an optimal reserve price in the first stage, it would also be revenue-maximizing. However, this assumes that the asset can effectively be allocated via an auction with a fixed number of bidders with no entry costs and considering this case would clearly go against the spirit of our analysis. Of course, if we did consider it, it would make the SASFE, where several potential buyers may incur the entry cost, appear even less optimal.

²⁶ As noted previously, a player should bid her value to being the marginal entrant. But, under NS, the marginal entrant is just as likely to win as any inframarginal entrant.

the fact that, in equilibrium, an early entrant with a relatively low value may be able to deter future entry in a very particular type of way. We will show that, with any degree of selective entry, entry-detering strategies change in an important way that can raise revenues significantly.

BK's sequential auction procedure works as follows.³¹ The seller approaches players in a random order. The first player decides whether to enter, incurring K if she does so, and, having found out her value, she can place a jump bid. If she does not do so, the standing price is either zero or a reserve set by the seller. The next player, who observes what happened in the first round, then decides whether to enter (incurring K), and, if he does so and the first player also entered, the two incumbents compete in a knockout button auction. The loser drops out and cannot subsequently re-enter the bidding. If the second bidder wins the knockout then, if he wants, he can place a new jump bid above the standing price at the end of the knockout. If the first player did not enter, then if the second bidder enters, there is no knockout but the second bidder can place a jump bid. The procedure is repeated until all potential bidders have had an opportunity to enter, at which point the object is allocated to the incumbent bidder at the standing bid.

Once we allow players to have signals before they enter, the only difference in how the game is played is that now a potential bidder's signal can also affect her entry decision. This difference, however, has a significant effect on equilibrium strategies. With no selection, BK show that there is excess deterrence from the perspective of the seller, which arises from the fact that jump bidding strategies involve a "semi-pooling equilibrium". To illustrate, suppose that the $(N - 1)^{\text{th}}$ bidder enters, and the previous incumbent exits the knockout at a price of b^s . BK show that the new incumbent's jump bidding strategy will be:

- place no jump bid if the new incumbent's value is less than V' ; or,
- place a jump bid equal to $b' > b^s$ if the new incumbent's value is greater than or equal to V' .

where V' is determined by the condition that, if the final potential bidder knows that the incumbent's value is at least V' , then it will choose not to enter.³² while b' is determined by the condition that an incumbent with value exactly V' should be indifferent between not placing any jump bid, in which case the next potential entrant will enter, and placing the jump bid b' and deterring entry. Subsequent potential entrants enter if and only if a jump bid of less than b' is placed.

Using the parameters assumed above, the bold line in the left-hand section of Fig. 3 shows the incumbent's jump bidding strategy when $b^s = 75$. In this case, $V' = 83.77$ and $b' = 81.43$, and the maximum revenues that the seller can realize are 83.77.³³ Note that if the next mover knew that the incumbent had a value of $V' + \epsilon$, where ϵ is small, entry would be profitable, but it chooses not to do so because it is unable to distinguish an incumbent with this value from ones with possibly much higher values.

In contrast, when entry is selective, the unique equilibrium jump-bidding strategies, under the D1 refinement (Cho and Kreps, 1987; Ramey, 1996), are fully separating for incumbent values up to $\bar{V} - K$.³⁴ Strategies change because, in our example, for any such incumbent value, the final player may enter if it receives a signal that is high enough. This gives an incumbent with a high value an incentive to distinguish itself from incumbents with lower values (because, by doing

³¹ We note that BK allow for the possibility that the number of potential bidders is stochastic, so that it is not known for sure at the beginning of the game. Our fixed N assumption is a special case of their framework.

³² Given that the $(N - 1)^{\text{th}}$ bidder did not receive a signal when deciding whether to enter, the pdf of the N^{th} bidder's belief about the new incumbent's value given that this value must be above b^s , is simply $\frac{F'(v)}{1 - F'(b^s)}$ for all $v \geq b^s$.

³³ These revenues would be realized when the incumbent has a value slightly less than 83.77, no jump bid is placed, entry occurs and the subsequent knockout ends at the incumbent's value. If the incumbent's value is more than 83.77, then the revenue will be 81.43.

³⁴ Incumbents with values above $\bar{V} - K$ continue to pool, but, because we are assuming \bar{V} is high, it is very unlikely that an incumbent will have a value this high (the probability is less than 10^{-4} for our parameters).

so, it will be able to deter more entry), and, under the refinement, this eliminates pooling equilibria (for values less than $\bar{V} - K$). Roberts and Sweeting (2013) show how the equilibrium jump bidding strategies in each round can be characterized by round-specific differential equations with lower boundary conditions where a new incumbent with a bid equal to the standing bid does not raise its bid.

The remaining lines in the left-hand section of Fig. 3 show the equilibrium bid schedule for several different values of α , including 0.95, 0.99 and 0.9975. The right-hand section of the figure shows the corresponding probabilities of entry as a function of the incumbent's value. When signals are very uninformative (α greater than 0.95), we see that only incumbents with values significantly above V' (the value that deters entry in the NS case) deter entry with high probability. With full separation, the entry decision of the final player will be socially optimal in the sense that entry will take place if and only if the expected increase in the value of the winner that occurs with entry is greater than the entry cost.³⁵

Given these bid functions, the seller's revenues will increase discontinuously when any degree of selection is introduced into the model because there is more entry and because the jump bids that high value incumbents place, which do tend to deter entry, are significantly above b' . As the degree of selection increases, incumbents with low values will deter some entry, which will not happen under NS, but because only potential entrants with quite low values are likely to be deterred, this does not tend to have a large effect on the seller's revenues. In our computations below we show how, during the course of the game with five rounds, this change in strategies can significantly increase the seller's revenues relative to the NS case, and also relative to the SASFE for a given degree of selection.

Roberts and Sweeting (2013) examine what happens under a wider range of parameters. When K is small, the difference in the values of the incumbents that deter entry in the sequential bidding auction under NS and with small degrees of selection, such as $\alpha = 0.99$, tends to fall slightly, while the difference in the level of deterring bid submitted by an incumbent with a very high value also becomes smaller (and, for very small K , can actually be higher in the NS case). This is illustrated in Fig. 4, which repeats the left-hand panel in Fig. 3, but for $K = 3$, $K = 1$ and $K = 0.5$. As a result of both of these changes, the change in expected revenues and surplus of the sequential bidding auction when selection is introduced (relative to NS) in the sequential bidding auction is much smaller, even though these is still a discontinuity in strategies.³⁶

4. Efficiency comparisons

Having introduced the various mechanisms, we now compare their performance. The standard measures of performance are efficiency (i.e., total surplus, defined as the value of winner less total entry costs) and seller revenues. We begin by considering total surplus, and then turn to the question of how surplus is split between the seller and the potential buyers. Our analysis uses the set of parameters introduced previously. When considering efficiency, we assume that there is no

³⁵ Of course, there are some inefficiencies associated with the procedure as a whole. In particular, those players who get to move first are more likely to win the object. However, when entry is at least moderately selective the difference in winning probabilities is small, and early-movers do not necessarily have higher profits, so that players' incentives to expend resources to try to get selected to move first may be limited. For instance, consider our example parameters when $\alpha = 0.6$. The first player wins with probability 0.2102 and the last player with probability 0.1857. The last player's expected payoff is actually higher than that of the first-mover (3.7362, compared to 2.2511), because the last player only enters when it is very likely to win at a relatively low price.

³⁶ As BK showed, the SASFE outperforms the sequential bidding auction under NS. The two changes discussed in the text lead to Roberts and Sweeting finding that the SASFE yields higher expected revenues than the sequential bidding auction when K is small and α is high (see Fig. 2 in Roberts and Sweeting, 2013, where the performance of the SASFE is improved by assuming that an optimal reserve price is used in that mechanism but not the sequential bidding auction). However, in these cases it is also true that the revenue advantage of the SASFE is small (typically less than 1%). In all cases, the sequential bidding auction dominates the SASFE in terms of efficiency.

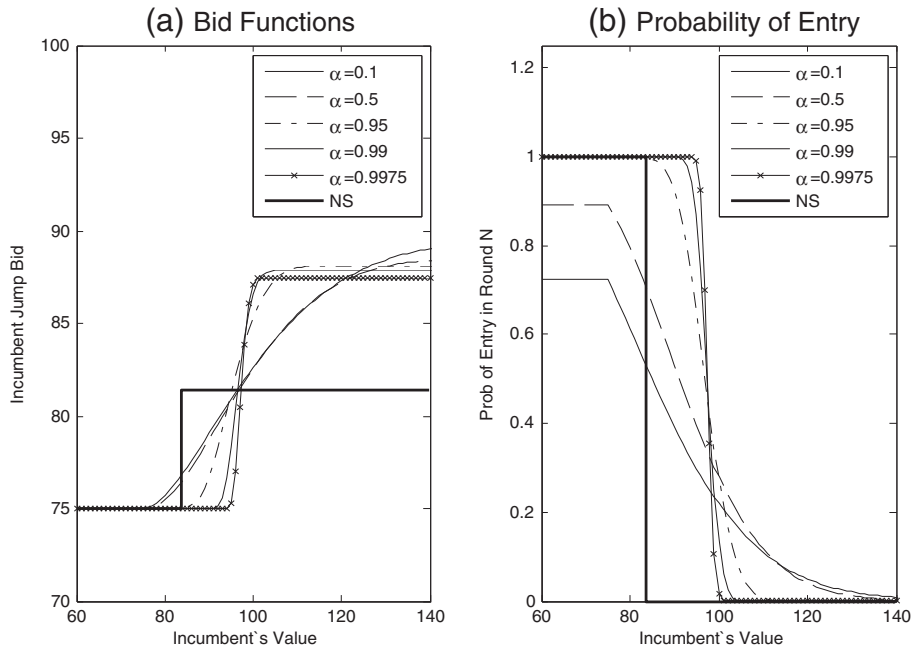


Fig. 3. Jump bid functions and entry probabilities as a function of the degree of selection in a sequential bidding auction model. Note: Diagram shows bid functions and entry probabilities in the penultimate round of a sequential bidding auction where the standing bid at the end of the previous knockout is 75 (so the incumbent's value must be at least 75).

reserve price in any mechanism, deferring consideration of reserve prices to the next section. We also assume that the seller sets a fixed n in the ERA in advance of the first stage to maximize expected revenues, although, with our parameters, it would make the same choices if it was aiming to maximize total surplus.

Fig. 5 reports the total surplus (5(a)) for the standard auction and the three alternatives, and a decomposition into the expected value of the winner (5(b)) and expected total entry costs (5(c)). We compute outcomes for the NS and FS models, as well as values of α between 0.01 and 0.99. Note that the values of α considered are not evenly spaced, with more values close to the polar cases, so that we can see

how slight deviations from the polar assumptions affect outcomes. It is immediately clear that small deviations from the NS assumption can affect outcomes more than small deviations from the FS assumption. For the ERA and sequential bidding auction, the NS cases are marked by discrete points as a reminder of the discontinuity of either the mechanism (ERA) or equilibrium strategies (sequential bidding auction) when one moves from a model with any degree of selective entry to NS.

Across all of the mechanisms, surplus increases when entry is more selective. The direction of this effect is expected as, when signals are more informative, it is more likely that it will be players with high values that have high signals, and so will enter or be chosen to enter.

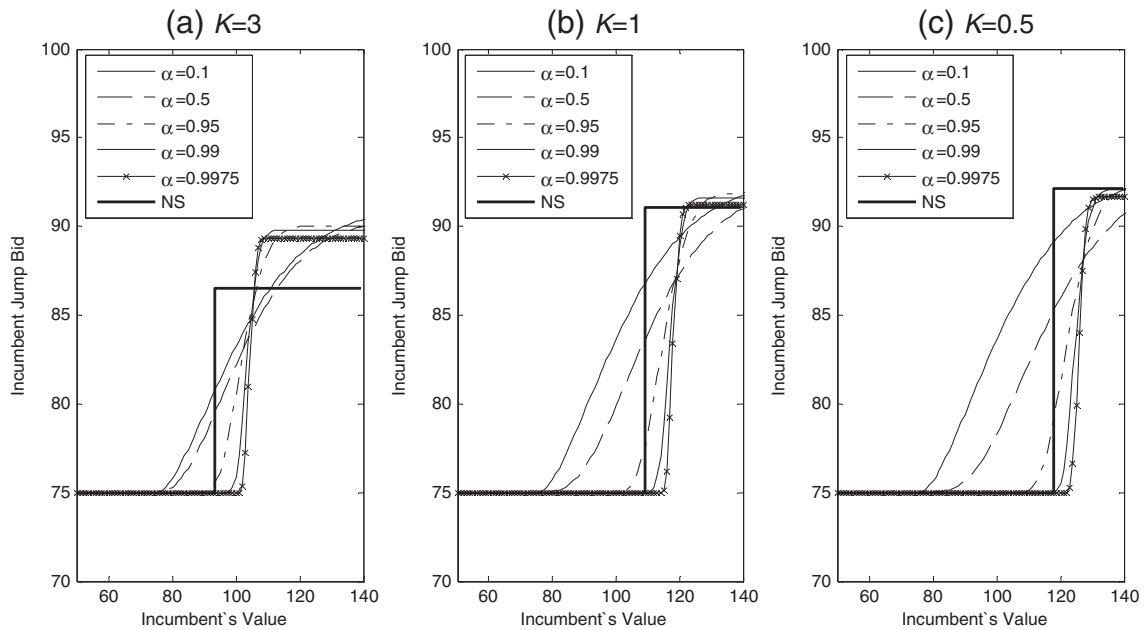


Fig. 4. Bid functions for three levels of K . Note: Diagram shows bid functions and entry probabilities in the penultimate round of a sequential bidding auction where the standing bid at the end of the previous knockout is 75 (so the incumbent's value must be at least 75).

Under NS, both the ERA and the sequential entry auction randomly select three players to enter. Excepting this case, the ranking of the mechanisms does not depend on the degree of selection, with the sequential bidding auction generating the highest expected surplus, followed by the ERA, the sequential entry auction and, finally, the SASFE.³⁷ However, it is also clear that the range of expected surpluses becomes much larger when one moves from the NS to the selective entry model (reflecting the fact that deterrence in the sequential bidding auction becomes more efficient), and tends to increase as entry becomes more selective.

Fig. 5(b) and (c) help to identify where the differences in surplus come from. As the degree of selection increases, surplus rises both because the value of the winner increases and the amount of entry falls, although, with the assumed entry cost, changes in entry tend to have a larger effect on surplus. For a given degree of selection, the SASFE and the two sequential auctions all lead to quite similar expected values of the winner, and for values of $\alpha \geq 0.7$, the value of the winner is actually highest in the SASFE even though the probability that the object is not sold to any player (so that the value of the winner is zero) is highest in this mechanism because of the uncoordinated nature of entry. The ERA attains the maximum possible value of the winner, i.e., it equals the expected highest value in the population of players, under the FS assumption because in this case the player with the highest value will definitely be selected to enter, because she has the highest signal, and will be allocated the object. On the other hand, when α lies between 0.8 and 0.95, the expected value of the winner is lower under the ERA because only two players are chosen to enter, and the ones selected may well not be the ones with the highest values. Note, however, that the problem is not (just) that there is less entry. As shown in panel (c), the expected amount of entry in the ERA is similar to that in the sequential bidding auction for these values of α . Instead the source of the relatively poor performance of the ERA in these cases is that it does not allow for there to be additional entry when the entrants turn out to have low values, whereas, when incumbents have low values, more players will tend to enter the sequential bidding auction.

The panels also reveal that the two sequential auctions produce very similar expected values of the winner, but that the sequential bidding auction achieves this with significantly less entry.³⁸ The key driver of the difference is that in the sequential entry auction it is very likely that the players who will get to move first will enter, even if they have low signals because they know that their entry decision will likely deter later players from entering, whereas in the sequential bidding auction, early players know that, in the fully separating equilibrium, they will only be able to deter entry if their values are high enough. To illustrate, suppose that $\alpha = 0.05$, so that players are quite well-informed about their values when they take their entry decisions. In the sequential entry auction, the first player to make an entry decision enters with probability 0.9652. In the sequential bidding auction, where its value will be revealed, it only enters with probability 0.3076. However, the probability that the first player wins in the sequential entry auction when it enters is only 0.2592, compared with 0.6619 in the sequential bidding case,³⁹ reflecting the fact that early-mover entry is less efficient when only entry decisions are sequential.

³⁷ In the case of the ERA this conclusion is not necessarily general, because of the restriction that n must be an integer greater than or equal to two. For example, when entry costs are high enough, expected entry into the SASFE may be less than 2, and the inefficiency associated with requiring two players to incur the entry cost may become more significant when entry is more selective.

³⁸ For example, on average, 2.16 players enter the sequential entry auction under the FS assumption, compared with 1.31 players for the sequential bidding auction. When $\alpha = 0.99$, the numbers are 2.87 and 2.44 respectively.

³⁹ Therefore, overall the probability that the first player wins in the sequential entry auction is 0.2502, compared with 0.2036 in the sequential bidding mechanism. For the remaining players in the order the probabilities of winning are [0.2303, 0.1981, 0.1715, 0.1498] in the sequential entry auction and [0.2036, 0.2012, 0.2003, 0.1973, 0.1976] in the sequential bidding auction. The similarity between these winning probabilities explains why the expected values of the winners are approximately the same. Note that under NS the first player enters with probability one in both of the sequential mechanisms.

In many settings, the seller may be able to increase the number of potential entrants by designing the object for sale appropriately or engaging in marketing activities. We therefore also look at how the expected total surplus changes when an additional player is added. The fact that there is an additional player will automatically tend to increase the expected highest value in the population of players (for our parameters, by 2.34), but the increase in expected surplus will depend on how likely this player is to enter and the change in total entry costs. As noted in footnote 3.2.1, we encountered some multiple equilibria in the sequential entry game when $N = 6$, so we ignore this mechanism in our comparisons (the equilibria in the other mechanisms are unique). The results are shown in the left-hand columns of Table 3.

As the degree of selection increases, the surplus gain from adding a player is larger for each of the mechanisms, reflecting the pattern that a player with a high value is more likely to enter, or be selected to enter, when there is more selection, and the fact that the increase in the expected number of entrants is smaller. In both of the sequential auctions and the SASFE, the average amount of entry also rises, whereas in the ERA the optimal n remains fixed. However, the change in entry has different effects in the SASFE and the sequential auctions. In the SASFE, because entry decisions are simultaneous and not coordinated, the equilibrium probability that no players enter (so that there is no surplus) rises even though the expected number of entrants increases. As a result, total surplus can actually fall when there are more players, and this happens in the SASFE when there is very little or no selection. On the other hand, surplus tends to increase in the sequential auctions, where some entry is effectively certain, for all values of α and, when selection is weak, the incremental surplus in the sequential auction is actually greater than in the ERA, which is a further reflection of the fact that fixing the number of entrants and selecting entrants based on their signals, as the ERA does, is particularly inefficient when signals are fairly uninformative.

5. Revenue comparisons

As already illustrated in the case of the SASFE (Table 1), selection affects the distribution of surplus between the seller and the potential buyers. For our mechanism comparison, Fig. 6 shows the seller's expected revenue (panel (a)) and the expected total bidder payoffs (panel (b)). In each mechanism, selection reduces expected revenues (excepting the move from NS to $\alpha < 1$ for the sequential bidding auction) and increases bidder payoffs, reflecting the fact that informative private signals create information rents.

Under the polar NS assumption, the ERA generates the highest expected revenue, with a small advantage over the SASFE and the sequential entry auction. Consistent with BK's theoretical result, the sequential bidding auction generates substantially lower revenue. Given our parameters, however, the ranking changes as soon as we introduce any degree of selection. The ERA and the sequential bidding auction generate higher revenues than the other mechanisms, and the advantage of the best alternative to the SASFE becomes significantly larger as the degree of selection increases. Of course, these rankings may depend on our selection of parameters, but we have found them to hold quite broadly when K is large enough so that it is likely that one or more bidders will not enter the SASFE when there is no selection (see footnote 3.2.2 for a discussion of why the ranking of the SASFE and the sequential bidding auction can change when K is small).

The relative performance of the sequential bidding auction and the ERA depends on the exact value of α , with the sequential bidding auction doing better for $\alpha = 0.9, 0.95$ or 0.99 , and the ERA doing relatively better with more selection. In the ERA, the relatively slow decline in expected revenues as α falls reflects two changes that tend to offset each other. Revenues in the second stage auction equal the second highest value of the selected entrants. Holding n fixed, this tends to increase as α falls because players' signals, which are used to select the entrants, become more informative. On the other hand, revenues from the first-stage auction tend to decrease as players with low signals know that their payoffs from

entering the auction, even if they are selected, are likely to be low. As a specific example, when $\alpha = 0.9$, expected first-stage revenues in the ERA are 15.74 (or 17.1% of total revenues), falling to almost zero under the FS assumption. Revenues in the sequential bidding auction fall when there is more selection as later entrants will become less inclined to enter unless their signal indicates that their value should be significantly above the value of the current incumbent. This reduction in entry, while socially efficient, tends to reduce the seller's revenues. The fact that which of these mechanisms performs best is sensitive to the degree of selection, and the fact that the advantage of the sequential bidding auction over the SASFE is also a non-monotonic function of α around 0.7, serve to illustrate the value of being able to identify the exact degree of selection in the entry process when choosing an auction design.

The sequential bidding auction always generate higher bidder payoffs, but the differences to the other considered mechanisms are

greatest in the polar NS and FS cases, and the advantage is quite small when there is selective entry and $\alpha \geq 0.5$. Even though an incumbent may be able to deter entry in these cases, it can be quite costly to do so, and, when signals are imprecise, players may enter the auction and subsequently find out that their values are not high enough to allow them to win the object. For a given value of α , the other mechanisms always generate quite similar payoffs for bidders. From a design perspective, the fact that a mechanism, such as the sequential bidding auction, is able to generate higher revenues for potential bidders, as well as the seller, may be important because, in the long-run, this type of design is likely to encourage more players to make the investments that will allow them to participate in auctions, benefiting the seller in terms of increased competition.

The right-hand side of Table 3 reports the change in expected revenue when another player is added. Bulow and Klemperer (1996) show that

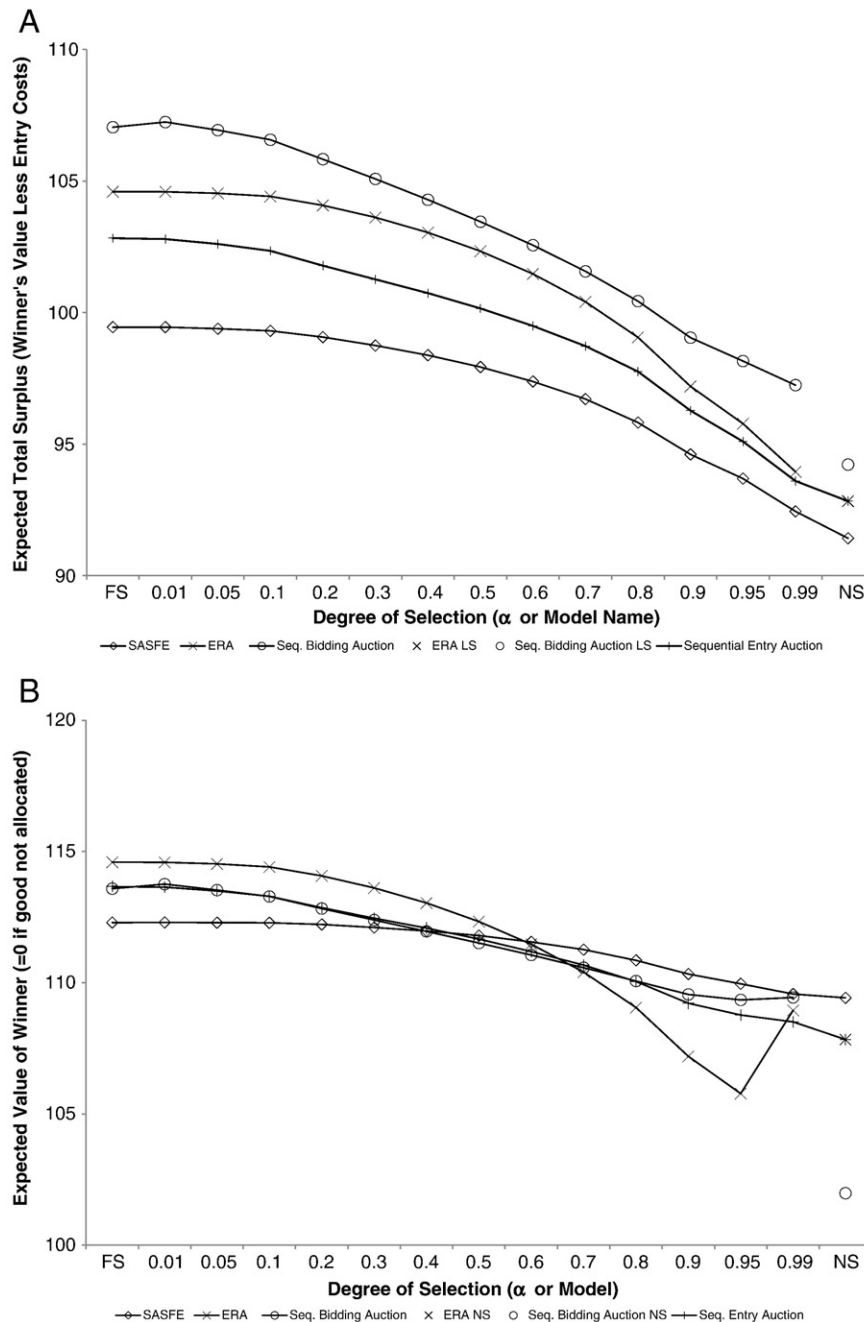


Fig. 5. (a) Expected total surplus. (b) Expected value of the winner. (c) Expected entry costs.




Fig. 5 (continued).

when the number of bidders is fixed, the seller always benefits more from adding an additional bidder than by using the optimal auction design. Under the NS assumption, because the potential buyers make no information rents in the SASFE or the ERA, the seller experiences the full efficiency increase or loss from adding a potential entrant, and the seller's expected revenues actually fall in an SASFE when an additional potential bidder is added. When entry is selective, the seller's revenues increase by more than the increase in total surplus, reflecting the fact that increased competition tends to reduce bidder payoffs. Interestingly, however, the increase in revenues when additional bidders are added is greater for the ERA and, especially, the sequential bidding auction. This suggests that the gains from using mechanisms other than the SASFE will become larger when more buyers are potentially interested in the object.

An obvious question is whether the revenue performance of these mechanisms can be significantly improved by introducing additional features, such as reserve prices. Fig. 7 shows expected revenues from the SASFE, with no reserve and a seller-optimal reserve (a seller-optimal entry fee leads to the essentially identical revenues as an optimal reserve), and the sequential bidding auction, with no reserve and a seller-optimal reserve that is common across rounds.⁴⁰

Under NS, the optimal reserve price for the SASFE is the value that the seller has of retaining the object, which is zero by assumption, so adding an optimal reserve has no effect on expected revenues. An optimal reserve can increase expected revenues in the sequential bidding auction under NS, although not by enough to reverse the SASFE's revenue advantage in this case.⁴¹ When entry is selective, a non-zero reserve

can be optimal in the SASFE, but they do not increase revenues by more than 0.4% for any value of α . In contrast, a strategic reserve in the sequential bidding auction can raise that mechanism's already larger revenues by as much as 4.7%, extending the revenue advantage of the sequential bidding auction over the SASFE to a really quite significant margin. Another way of framing these results is that, when entry is selective, the gain in revenues of switching from an SASFE to either an ERA or a sequential bidding auction can be many times greater than the value of setting an optimal reserve in the SASFE.

6. Asymmetric bidders, bid preferences and selection

In this section we examine how selective entry impacts how bid preference programs affect revenues and efficiency in a setting where players are asymmetric. Bid preferences are widely used, partly to meet distributional targets, such as allocating a certain proportion of

Table 3 Changes in Outcomes when the Number of Players is Increased from Five to Six.

α	Expected Total Surplus			Expected Revenues		
	SASFE	ERA	Sequential Bidding	SASFE	ERA	Sequential Bidding
0 (FS)	1.58	2.34	2.29	2.14	2.92	4.26
0.01	1.56	2.33	2.25	2.07	2.92	4.05
0.05	1.55	2.32	2.20	2.06	2.91	3.97
0.1	1.51	2.28	2.14	2.05	2.88	3.86
0.2	1.47	2.20	2.05	2.03	2.83	3.66
0.3	1.41	2.08	1.94	1.98	2.73	3.45
0.4	1.32	1.95	1.85	1.86	2.62	3.14
0.5	1.21	1.80	1.73	1.77	2.45	2.87
0.6	1.09	1.63	1.58	1.67	2.24	2.55
0.7	0.87	1.43	1.42	1.47	1.99	2.26
0.8	0.70	1.17	1.23	1.22	1.66	1.93
0.9	0.37	0.80	0.95	0.78	1.21	1.55
0.95	0.17	0.57	0.79	0.56	0.84	1.33
0.99	-0.19	0.32	0.49	0.03	0.38	1.08
NS	-0.48	0	0.06	-0.48	0	0.15

Notes: Expected outcomes based on 500,000 simulations.

⁴⁰ In the sequential bidding auction, we model the reserve price as being an initial standing bid chosen by the seller. It should be possible to do even better by allowing the seller to change this reserve price across rounds until entry occurs, but for simplicity we do not try to calculate the gains from using this type of policy. Including a positive reserve price in the all-pay first-round of an ERA could lead to some players refusing to participate, while it is not necessarily attractive to think that the seller would be willing to subsidize second-round entry costs if the bidder subsequently refused to meet a second-round reserve price. Computing the optimal reserve in a sequential entry auction would involve an excessive amount of computation.

⁴¹ Specifically, it is optimal for the seller to set a reserve price above V' (for our parameters a reserve of 92.63 is optimal). For this reserve price, entry will occur until a buyer indicates that they are willing to meet this reserve, in which case entry stops, and the object is sold at the reserve. Therefore, the mechanism essentially becomes one where the seller sequentially offers the good to each potential buyer at a fixed price.

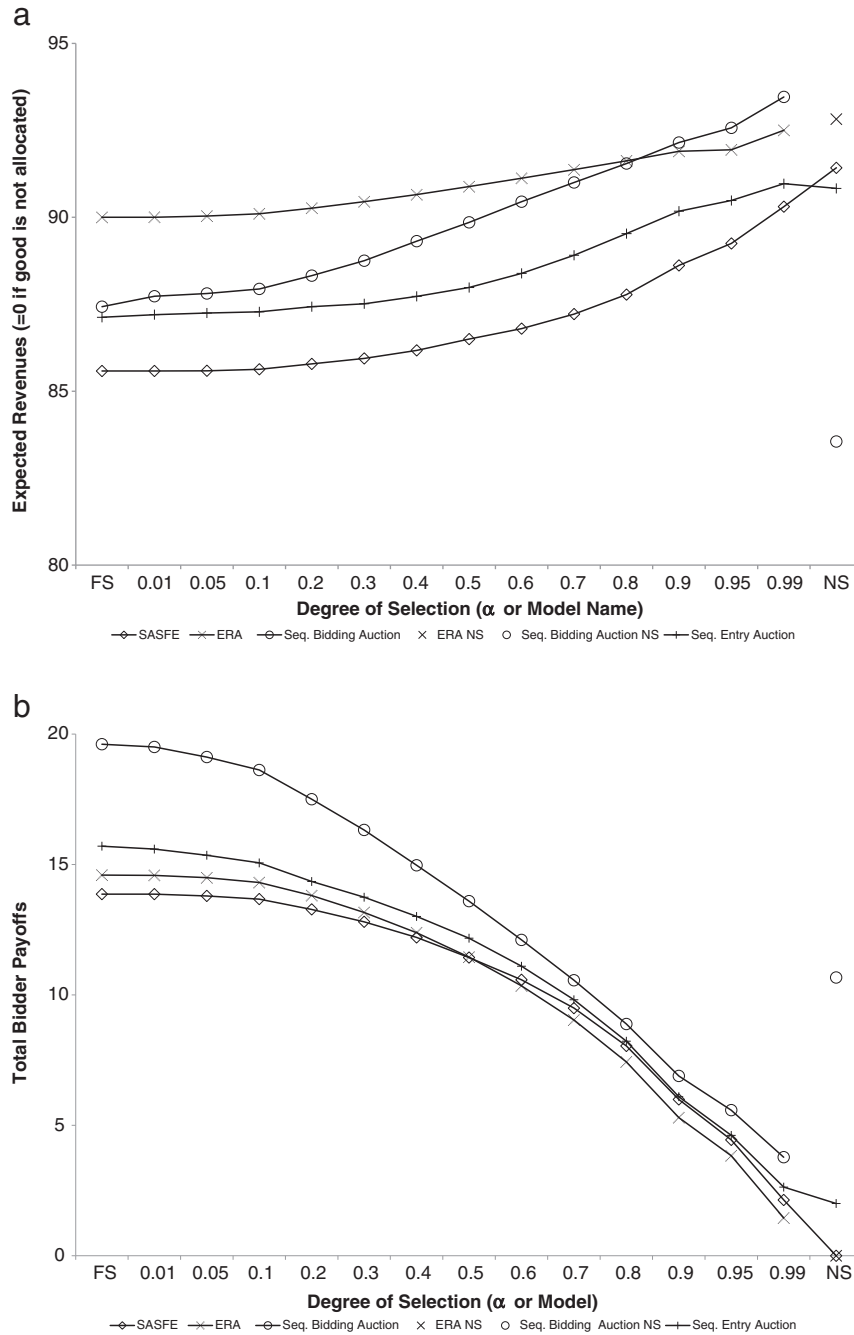


Fig. 6. (a) Expected revenues. (b) Expected sum of bidder payoffs.

contracts to small businesses or companies owned by minorities,⁴² but they may also increase seller revenues when one group of players tends to have lower values. In this section, we use a model where players are asymmetric and examine how selection affects the level of bid preference that a seller that wants to maximize its revenues should choose.

In the existing literature, bid preference programs have been analyzed primarily under NS. For example, Athey et al. (2013) (ACL) use an estimated NS model to show that bid preferences may be preferable to other distributional schemes, such as set-asides, in the context of US Forest Service timber auctions, and they estimate that preferences of

more than 20% would maximize the revenues of the agency. Krasnokutskaya and Seim (2011) analyze bid preferences in state highway procurement auctions using an NS model with heterogeneous entry costs. Our analysis will be closer to that of ACL in that, like them, we will keep the entry decisions of stronger bidders fixed when we introduce a bid preference.⁴³ Hubbard and Paarsch (2009) analyze the effects of bid preferences under the FS model assumption that players know their values, although they assume that players are symmetric

⁴² For example, the FCC grants bid credits to small businesses in some spectrum auctions (Congressional Budget Office, 2005), while many states give credits to small or minority owned businesses in procurement auctions.

⁴³ In their setting, Krasnokutskaya and Seim predict that granting weaker players bid preferences can increase procurement costs because the effect on the entry decisions of stronger bidders can be large. However, allowing for selection could also change this conclusion because the strong bidders that would cease to enter would tend to have higher costs than those that remain. If there are multiple equilibria in the entry game, as is quite possible when the entry decisions of both types are endogenous, the effectiveness of bid preferences may also depend on the equilibrium that is played.

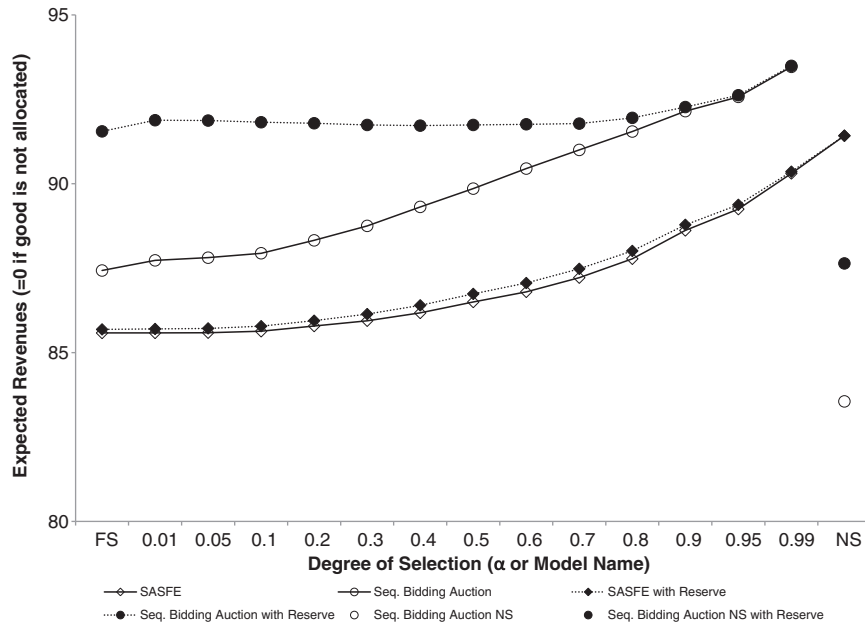


Fig. 7. Expected revenues from a sasfe and a sequential bidding auction with seller-optimal reserve prices.

when preferences are applied. The approach that we use to solve our auction model with asymmetries, bid preferences and selective entry extends their approach to a more general setting.

We illustrate the effect of bid preferences by extending the model considered above to allow for two types of bidders. In our example, there are exactly $N_s = 2$ ‘strong’ (s) bidders, with $\ln(V_i^s) \sim N(4.5, 0.2)$, and $N_w = 4$ ‘weak’ (w) players, with $\ln(V_i^w) \sim N(4.4, 0.2)$, with both distributions truncated to $[40, 190]$.⁴⁴ With these parameters, the average value of a strong bidder is 91.83, the average value of a weak player is 83.06 and the probability that one of the weak players has the highest value is 0.49.

We assume that the strong bidders and the entering weak bidders compete in a first-price sealed bid auction for the object, where they submit bids uncertain about how many (other) weak bidders may have entered the auction. The bid preference affects the allocation of the good, but does not affect the payments that are made. For example, if the proportional bid preference is ρ (e.g., 0.1 or 10%), the auctioneer allocates the object treating the bids of weak bidders as if they were multiplied by $(1 + \rho)$, while using the actual bids of strong bidders. On the other hand, if a weak bidder wins the good based on these inflated bids, it only pays its actual bid.⁴⁵

We choose to assume that the strong bidders enter for sure, and that only the entry decisions of weak bidders, who receive signals before deciding whether to incur an entry cost, $K = 5$, are endogenous. We do so partly to compare our results to ACL, but also to side-step the issue that there may be multiple type-symmetric equilibria in the entry game where the entry of both types is endogenous.⁴⁶

As in the SASFE considered above, weak bidders will enter if their signals are above a threshold S_w^* , and then in the auction, bidders of

⁴⁴ For numerical reasons it is convenient to have a lower truncation point that is above zero. The probability that a value drawn from an untruncated distribution would be less than 40 or greater than 190 is very small for both types.

⁴⁵ Note that ACL actually consider a bid subsidy scheme where actual bids are used to determine the allocation of the good, but favored bidders only pay some proportion, $\frac{1}{1+\rho}$, of their bids. While bids will be different, the allocation and payments should be identical to our model when we consider a bid preference of ρ .

⁴⁶ Roberts and Sweeting (2013) consider a setting with two types and make the equilibrium selection assumption that the strong type has the lower entry threshold, on the basis that there is always exactly one equilibrium of this type and the implication that stronger types will certainly be more likely to enter is attractive. However, when the weak players receive bid preferences, the analysis would become more complicated, and the uniqueness property might change as the size of the bid preference increases.

both types will submit bids according to type-specific bid functions. We assume that bidders place bids uncertain about the number of weak players that have entered. The Bayesian Nash equilibrium consists of this threshold and type-symmetric bid functions for each type. Define $H^w(v)$ and $H^s(v)$ as the probabilities that a player of a particular type (w or s) either does not enter or enters and has a value less than v .⁴⁷ $H^w(v)$ will therefore depend on the threshold. The probability that a weak bidder will win with a bid b_w is

$$[H^s(\beta_s^{*-1}(b_w(1 + \rho)))]^{N_s} [H^w(\beta_w^{*-1}(b_w))]^{N_w-1}$$

where $\beta_w^*(v)$ and $\beta_s^*(v)$ are the equilibrium bid functions of each type of bidder and the $b_w(1 + \rho)$ term reflects the fact that a strong bidder can only win if he bids at least $b_w(1 + \rho)$. The equilibrium weak bidder bid function $\beta_w^*(\cdot)$ will then be determined by the solution to the optimization problem

$$\beta_w^*(v) \equiv \arg \max_b (v - b) [H^s(\beta_s^{*-1}(b(1 + \rho)))]^{N_s} [H^w(\beta_w^{*-1}(b))]^{N_w-1}.$$

The first order condition associated with this optimization problem gives a differential equation

$$1 - (N_w - 1)\beta_w^{*-1}(b)(\beta_w^{*-1}(b) - b) \left[\frac{H'_w(\beta_w^{*-1}(b))}{H_w(\beta_w^{*-1}(b))} \right] - N_s \beta_s^{*-1}(b)(\beta_s^{*-1}(b) - b)(1 + \rho) \left[\frac{H'_s(\beta_s^{*-1}(b(1 + \rho)))}{H_s(\beta_s^{*-1}(b(1 + \rho)))} \right] = 0$$

and there is a lower boundary condition where

$$\beta_w^*(r) = r. \tag{1}$$

In our calculations we assume a type-independent reserve price of 50, although, because we assume that two strong bidders enter, this reserve almost never binds. There is also an endogenously determined upper boundary condition where $\beta_w^*(\bar{V}) = \bar{b}$. Similar equations define

⁴⁷ Of course, we are assuming that the strong type enters for sure.

Table 4
Value of K that Leads to an Average of 1.5 Weak Players Entering, with No Bid Preference, as a Function of α .

Degree of Selection (α)	Required Entry Cost (\hat{K})
0.01	0.523
0.05	0.575
0.1	0.655
0.2	0.847
0.3	1.066
0.4	1.297
0.5	1.536
0.6	1.776
0.7	2.017
0.8	2.255
0.9	2.484
0.95	2.591
0.99	2.669

the equilibrium bid function for strong bidders, except that the upper boundary condition will be that $\beta_s^*(\bar{V}) = (1 + \rho)\bar{b}$.

S_w^* is determined by the zero profit condition

$$\int_r^{\bar{b}} (\beta_w^{*-1}(b) - b) f_v^w(\beta_w^{*-1}(b) | S_w^*) [H^s(\beta_s^{*-1}(b(1 + \rho)))]^{N_s} [H^w(\beta_w^{*-1}(b))]^{N_w - 1} db = K, \tag{2}$$

where $f_v^w(\cdot | s)$ is the conditional density of a weak bidder's value, computed using Bayes' Rule, given she receives a signal s .

To illustrate the effect of selection on the optimal level of bid preference, we consider the following problem: suppose that a researcher has data from a set of identical auctions where there is no preference scheme, and he wants to know how a particular level of bid preference would change outcomes. We are interested in how his answer may depend on the assumed value of α . To do this exercise, we change the value of K so that, when we change α , we keep the expected number of weak bidders with no preference fixed.⁴⁸ This effectively involves setting S_w^* , solving the equilibrium bidding game, and then identifying the level of K that makes the zero-profit condition (Eq. (2)) hold. By doing so, we can also illustrate how changing α affects the level of K that a researcher would likely estimate.⁴⁹ However, this approach does create some differences to the results presented in previous sections. In particular, in the model with symmetric bidders and fixed K , seller revenues decreased with selection, as fewer players entered. Now that we are holding entry fixed, seller revenues tend to increase with selection as the entrants tend to have higher values.

To solve for the bid functions, we use the Mathematical Programming with Equilibrium Constraints (MPEC) approach (Su and Judd, 2012; Hubbard and Paarsch, 2009), where we use the AMPL programming language and the SNOPT solver. We express the inverse bid functions of each type as a linear combination of the first P Chebyshev polynomials (we use $P = 25$), scaled to the interval $[r, \bar{b}]$. When we are solving the model with no bid preference, S_w^* is set to get the right amount of entry by weak players, and the choice variables in our programming problem are, therefore, $2P$ Chebyshev coefficients, K and the value of the upper boundary condition for weak entrants (\bar{b}).

⁴⁸ One could also, of course, examine how outcomes change, as a function of the level of bid preference, keeping K fixed and allowing the amount of weak player entry to change with the degree of selection. In our example, when we hold K fixed, the amount of weak bidder entry with no bid preference can change quite dramatically with α , and this change has a large effect on outcomes. We view our formulation as a more natural way to illustrate how a researcher's conclusions about the optimal level of bid preference would be affected by his assumptions about selection in a setting where weak bidders already frequently participate in auctions.

⁴⁹ Of course, if the researcher was actually estimating the model, he would also be estimating the distribution of values, and these would also be affected by the assumed α . For our illustration, we assume that the researcher knows the true value distributions.

When we solve the model with a bid preference, we take K as fixed, but also solve for the equilibrium entry threshold, S_w^* .

Table 4 shows how the level of K, \hat{K} , required to rationalize why 1.5 weak bidders enter on average changes with α . In this part of the analysis we consider values of α from 0.01 to 0.99, without solving for the polar cases. Consistent with the logic outlined in footnote 2.2, weak players are more willing to enter when signals are less informative, so that as α increases, entry costs must also increase to explain why some of these players choose to remain out. The change in \hat{K} is large: moving from $\alpha = 0.01$ to $\alpha = 0.99$, increases the required K by a factor of 5, from 0.6% of a weak player's mean value to over 3.2%. This pattern may also explain why researchers typically estimate entry costs that seem implausibly high when they assume NS.⁵⁰

Given these values of K , we analyze how bid preferences impact auction outcomes, as a function of α . To do so, we consider a finite set of different levels of bid preferences, $\rho = \{0, 0.025, 0.05, 0.075, 0.1, 0.125, 0.15, 0.175, 0.2\}$. We then solve the model, including the equilibrium bid functions of each type, for each $(\alpha, \hat{K}(\alpha), \rho)$ combination. We then simulate 100,000 outcomes for each case in order to calculate expected outcomes.

Fig. 8 shows four expected outcome measures. Consistent with the logic of optimal auctions, which is developed in the context of models with exogenous entry, the seller's revenues (8(a)) can be increased using bid preferences. In addition to making existing weak entrants more likely to win the auction, which should cause strong bidders to shade their bids less, a bid preference will also encourage more weak players to enter. This entry effect will encourage both weak and strong bidders to bid more aggressively. Fig. 8(b) shows how bid preferences affect weak bidder entry for different degrees of selection, with larger effects as α increases.

The most striking finding is that the level of bid preference that is predicted to maximize the seller's expected revenues is much smaller when entry is assumed to be very selective. This suggests that any finding that large bid preferences are optimal when NS is assumed may be very sensitive to that assumption. For the lowest α that we consider (0.01), the revenue-maximizing preference is $\rho = 0.025$; for values between 0.05 and 0.5 the revenue-maximizing preference is $\rho = 0.05$; whereas, for $\alpha = 0.9$, for example, $\rho = 0.125$ is optimal. The fact that large bid preferences are optimal when the degree of selection is small reflects the fact that, in this case, a preference has a large effect on bidder entry (7(b)), and that the additional entrants that are drawn in will tend to have value distributions that are similar to those that would enter without a preference, so that they will tend to be equal competitors to the inframarginal entrants. In contrast, when entry is selective, the new entrants will tend to have relatively low values, and the effect on incumbent bidding will be more limited. As an illustration, Fig. 9 shows how the equilibrium bid functions of both types change when a bid preference of $\rho = 0.1$ is introduced for $\alpha = 0.1$ (9(a)) and $\alpha = 0.9$ (9(b)). Considering a strong bidder with a value of 120, the increase in the bid that follows from the bid preference is one percentage point larger when $\alpha = 0.9$ than when $\alpha = 0.1$ (5.2% vs. 4.1%). The decrease in the weak bidder's bid is approximately one-half-of-one percentage point smaller (2.1% vs. 2.6%). These differences naturally make the use of a larger bid preference, even though it may allocate the good to a

⁵⁰ For example, ACL estimate mean entry costs for a sample of US Forest Service timber auctions of \$7.54/mbf, compared to the Roberts and Sweeting (2013) estimate of \$2.05/mbf when they allow for selection (although these estimates are based on different samples of data) and a forester's estimate of entry costs of around \$3/mbf. Bajari et al. (2010) and Krasnokutskaya and Seim (2011) estimate that average entry costs into highway procurement auctions are 4.5% and 3% respectively of the engineer's estimate of the cost of completing the contract. These are also significantly higher than those of Bhattacharya et al. (2014) (1.5%), also in a highway procurement context, and most industry estimates (Park and Chapin, 1992; Halpin, 2005) that lie in the range 0.25–2%.

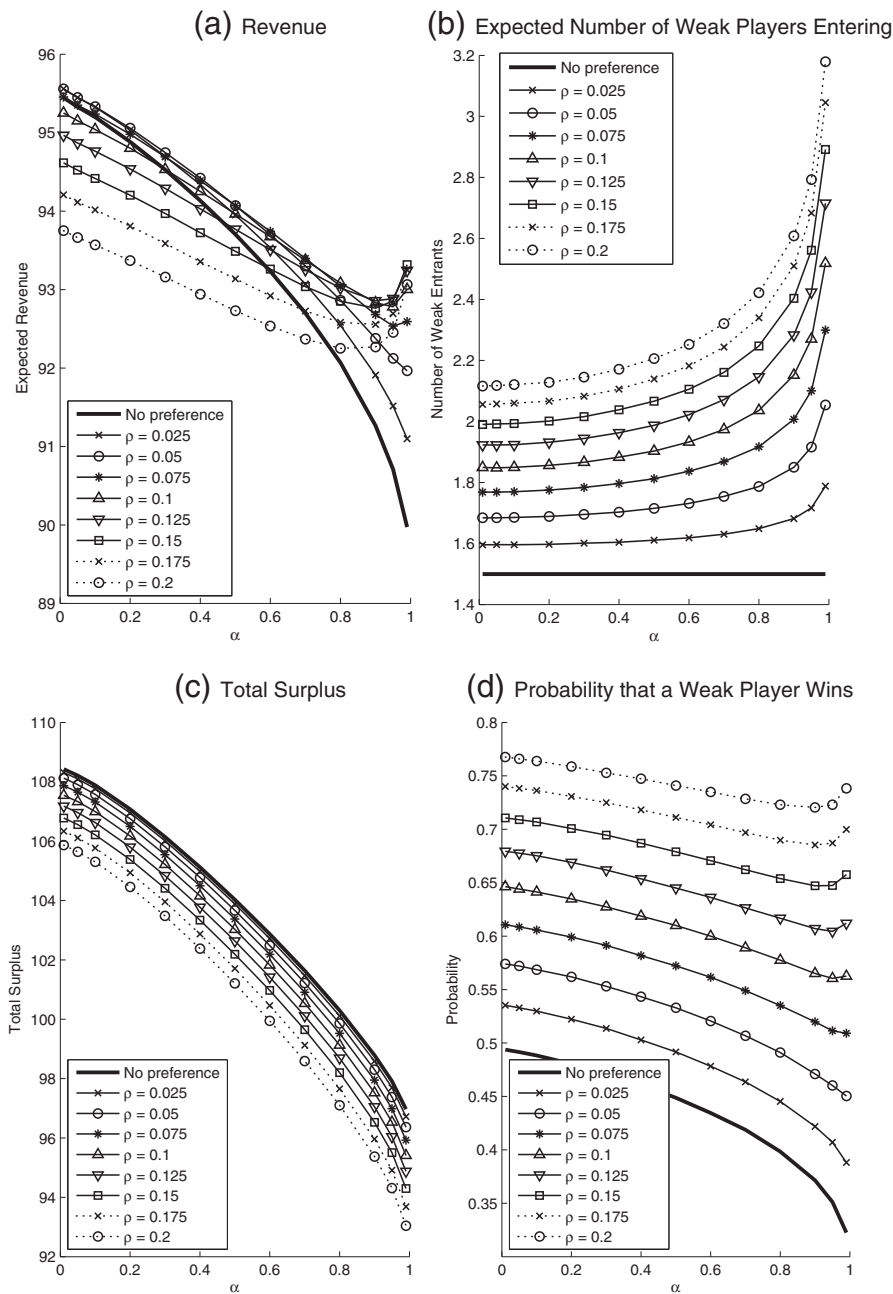


Fig. 8. Effect of bid preferences in a model with asymmetric bidders. Note: In these figures K varies with α so that the expected entry of weak bidders when there is no preference is constant.

bidder with a lower nominal bid, more attractive when there is less selection.

Fig. 8(c) shows that the effect that preferences have on total surplus is fairly independent of the assumed degree of selection. Even though preferences will make strong bidders shade their bids less and weak bidders shade their bids more, all of the preferences that we consider tend to increase the probability that the good is allocated to a weak entrant that does not have the highest value (8(d)), while simultaneously tending to increase entry costs. The former effect is larger when there is assumed to be more selection, which offsets the fact that the social cost of increased entry is lower when there is more selection because, to hold the level of weak bidder entry fixed, K is smaller.

As noted above, one motivation for using bid preferences is to achieve a distributional objective that weak players should win more often. While bid preferences do have this effect, with or without selection, it is also noticeable that when more selection is assumed, the

probability that a weaker player wins also rises.⁵¹ As a point of comparison, the probability that a weak player has the highest value is 0.49, which is only slightly greater than the probability that a weak player wins when α takes on the lowest value that we consider and there is no bid preference. This reflects the fact that, even though the expected amount of weaker bidder entry is (weakly) lower with selection given the way that we are changing K ,⁵² selection leads to the weaker players that are most able to win entering. This suggests that a seller with distributional objectives might try to achieve them by increasing selection (for example, by providing potential buyers with better

⁵¹ The one exception to this is for values of α above 0.95 when large bid preferences are in place.

⁵² When there is no preference, the amount of weak bidder entry is the same regardless of the degree of selection. When preferences are introduced, the amount of weak bidder entry increases by less when more selection is assumed.

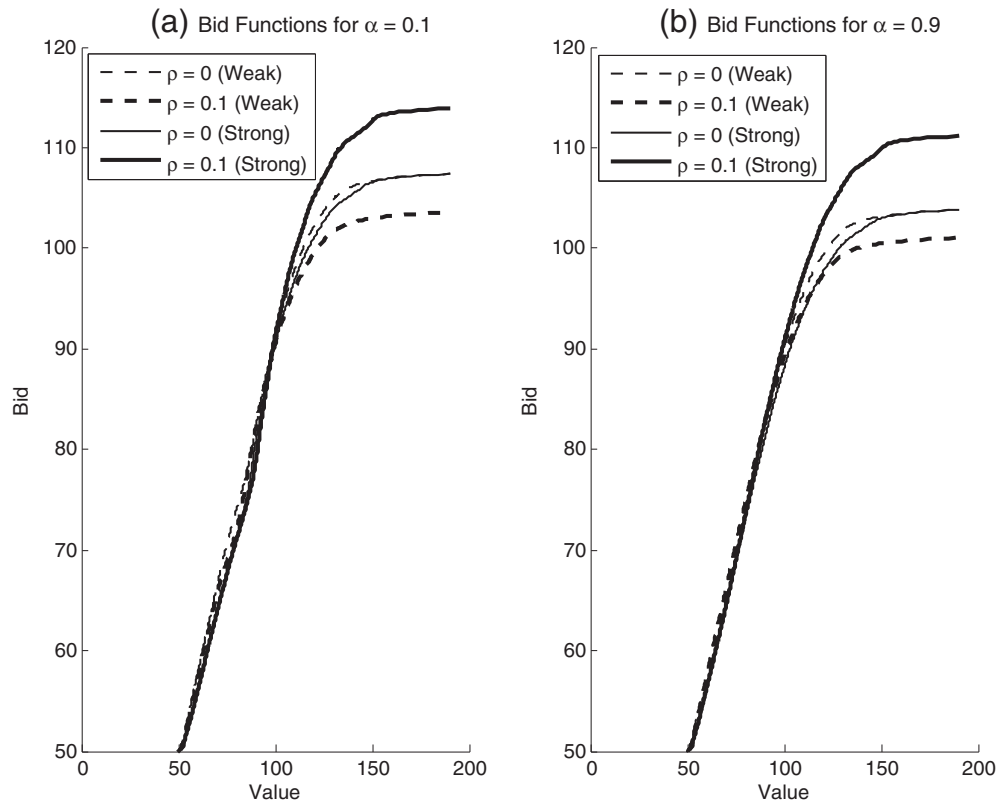


Fig. 9. Bid functions with and without a 10% bid preference.

information about the good being sold), rather than, or in addition to, introducing bid preferences.

7. Measuring selection

The previous results suggest that the degree of selection can significantly affect both the direction and the size of gains from changing an auction design. These results, however, would be of limited value if we cannot identify how selective the entry process is in different real-world contexts. Fortunately, there has been significant progress in the last few years when it comes to the identification and estimation of auction models with endogenous and selective entry. We highlight some of the most significant contributions in this section.

The primitives of our model that a researcher would like to estimate in order to do interesting counterfactuals, such as evaluating different auction designs, are the joint distribution of players' values and signals, and the level of entry costs. Assuming that players are symmetric, the data is assumed to contain the number of players, the number of entrants, and some information on the bids of entrants. The fundamental identification problems are that we never observe signals directly and that, with selection, the distribution of entrants' values, which should map into their bids, will be different from the distribution of values in the population. However, at least some of these problems occur in other economic settings where there is selection. The intuition from these settings is that we can potentially identify the population effects and the degree of selection when there is some exogenous source of variation in which, or how many, agents choose to enter. The same logic holds in the current setting, although there is an additional layer of complexity introduced by the fact that entry rules and bidding behavior should be determined by the equilibrium of a multi-agent game rather than by a single agent optimization problem.⁵³

⁵³ For example, when some exogenous variable affects the equilibrium amount of entry, it may also affect equilibrium bidding strategies in first-price auction contexts.

The focus of work on non-parametric identification and testing has been on models of first-price or second-price button auctions (where it may be possible to observe the bids of multiple entrants) and potential entrants are symmetric. Symmetry implies that variation in the amount of entry then comes from variation in auction-level variables, such as the number of potential entrants, the reserve price or the level of entry costs. [Marmer et al. \(2013\)](#) show that one can non-parametrically distinguish the NS, FS, and partially selective entry models using exogenous variation in the number of potential bidders and estimates of the quantiles of the value distributions of entrants conditional on the number of potential entrants (which can be calculated from inverting bid functions, which will be specific to the number of potential entrants, using the methodology proposed by [Guerre et al., 2000](#)).⁵⁴ The basic intuition is that under NS these quantiles should be invariant to the number of potential entrants, whereas with selection the quantiles should tend to increase with the number of potential entrants as the equilibrium signal threshold for entry should rise. Under FS, where bidders will only enter if their values are above the threshold, the shift in the lower quantiles of the distribution of values should be particularly sharp.

One can also derive other tests of the polar models from some of the results described above (see [Coey et al., 2014](#) for further examples). For example, under the NS assumption, expected revenues should fall as the number of potential entrants increases. While models with selection do not necessarily have a positive relationship, a rejection of a negative

⁵⁴ [Li and Zheng \(2009\)](#) compare the model fit of an NS model with a common entry cost, an NS model with heterogeneous entry costs and an FS model with a common entry cost in a procurement setting, whereas [Li and Zheng \(2011\)](#) compare the fit of the first and last models in the context of high-bid timber auctions. They find that the FS model provides a better fit in the timber setting, whereas they find that the common entry cost NS model fits best in the procurement setting.

relationship, based on exogenous variation in the number of potential entrants, could lead one to reject the NS model, as could an observation that there is some but incomplete entry by bidders of multiple types when potential bidders are asymmetric.⁵⁵

Gentry and Li (2014) consider, more explicitly, the problem of whether the joint distribution of values and signals is identified. They prove point identification when there is continuous and exogenous variation in the equilibrium entry threshold, for example due to an observable source of variation in the entry costs or exogenous variation in the reserve price. On the other hand, when there is only discrete variation, for example due to variation in the number of players, the joint distribution may only be bounded, but, of course, with enough variation these bounds may be quite tight.⁵⁶ They show that these results can be extended to settings where there is unobserved auction heterogeneity, which may also affect the observed number of potential entrants, as long as some exogenous variation in factors affecting entry remains. Of course, once the joint distribution of values and signals is identified, entry costs are identified from the fact that a player that receives the threshold signal should be indifferent to entering.

Roberts and Sweeting (2015), Roberts and Sweeting (2013) and Bhattacharya et al. (2014) currently provide the main empirical applications where models of selective entry are estimated. In all three cases a fully parametric approach is taken, partly because of the need to control for covariates, but also to address the fact that exogenous sources of variation in the entry thresholds are typically limited in practice. Roberts and Sweeting (2015) and Roberts and Sweeting (2013) estimate a model of second-price auctions using data from US Forest Service timber auctions in the Pacific Northwest. Like ACL and Athey et al. (2011), they allow for asymmetries between sawmills and logging companies, and for unobserved auction heterogeneity. From a modeling perspective, asymmetries between bidder types provide both a potential opportunity and a problem for identifying the parameters. The opportunity comes from the fact that variation in the number (or characteristics) of players of one type should lead to variation in the entry strategies of other types of players, providing a new source of identification. Indeed, as Gentry and Li (2014) note, when there is continuous variation in player types, as might be created for example by the distance of a potential bidder's depot to a construction site, as is often observed in highway paving auctions, or from a sawmill to the location of the tract in a timber auction (as in Li and Zhang, forthcoming), this can lead to point identification of the model. On the other hand, there may be multiple equilibria even when one restricts oneself to strategies that are symmetric within-type, requiring an equilibrium selection assumption (usually that only one equilibrium is played in the data given a set of observed auction and potential bidder characteristics) or an estimation approach to deal with the incompleteness of the model (Tamer, 2003) that multiplicity can generate. In an environment with two types that differ only in the location parameters of their value distributions, one can show that there will always be exactly one equilibrium where the stronger type has a lower entry threshold implying that they are more likely to enter. Assuming that this is the equilibrium that is played, Roberts and Sweeting (2015) and Roberts and Sweeting (2013) estimate a

random coefficients model of the structural parameters. They find that the mean value of α for the auctions in their sample is 0.6, indicating moderately selective entry and potentially significant gains from using non-standard designs.⁵⁷

Bhattacharya et al. (2014) estimate a parametric model of selective entry into low bid procurement auctions with symmetric bidders. They also estimate that entry is moderately selective (mean value of α is 0.5) and they use the approach of Gentzkow and Shapiro (2014) to illustrate which moments of the data parametrically identify the parameters. The results are broadly consistent with the intuition from Marmer et al. (2013) and Gentry and Li (2014) in that changes in the number of realized entrants as the number of potential entrants varies play an important role in determining the degree of selection. All of these papers take a full information approach to estimation in the sense that it is necessary to be able to solve the selective entry and bidding games as part of the estimation process. The results in Gentry and Li (2014) suggest that it may be feasible to use a two-stage approach to estimating selection models with real-world data. It would also be interesting to investigate how the Haile and Tamer (2003) inequality-based approach to estimating open-outcry auction models, which assumes a fixed number of bidders, could be applied to settings with endogenous and selective entry.

8. Conclusion

This article has argued that it is important to account for the selectivity, as well as the endogeneity, of entry when trying to evaluate different auction designs in real-world settings. In the particular example considered an increase in the degree of selection tends to increase both the efficiency and revenue gains from deviating from the 'standard auction with simultaneous and free entry' format, and it also tends to reduce the value, to the seller, of large bid preferences in settings where bidders are naturally asymmetric. Recent advances in the empirical literature make it feasible to estimate parametric models of selective entry, and to argue that the degree of selection is non-parametrically identified.

In our discussion we have largely treated the degree of selection as a fixed parameter, rather than some feature of the auction that the designer gets to choose. In practice, in many settings the seller makes choices about what information will be provided to potentially interested parties, and it may be able to hire third parties who could provide independent assessments of the object being auctioned, with the aim that bidders can make more informed choices. When entry costs are fixed, increasing the degree of selection tends to increase efficiency but reduce the seller's revenues by reducing entry/competition and increasing information rents. However, if providing better information also reduces entry costs, so that the amount of entry is maintained, as in Section 6, then the change may increase both efficiency and revenues. Understanding how auction design, including features that are often ignored in the literature, such as the information distributed to potential bidders, affects both entry costs and selection seems to us to be an important direction for future research. It would also be interesting to extend our framework to allow for the possibility that there is a common value element in players' valuations. In this case, it may be even more desirable for the seller to regulate entry, in order to reduce the effects of the winner's curse, but it is an open question as to what type of mechanisms might achieve this type of control in the most efficient way.

The discussion also has some relevance for how to think about entry in non-auction settings. A key feature of the selective entry model is that marginal and inframarginal entrants are not necessarily alike, whereas it is quite common in entry settings to assume either that entrants only find out any unobserved characteristics post-entry or that the unobserved characteristics that they do observe simply reflect differences

⁵⁵ In a type-symmetric equilibrium with multiple bidder types under NS, players of a given type must either all mix over entry, enter with certainty, or not enter for sure. However, with a common K , it will generically not be possible to have two types, with different value distributions, that are both willing to mix. Therefore observing some but incomplete entry by more than one type can be used to reject the NS model, although this conclusion is dependent on the assumption that only type-symmetric equilibria are played.

⁵⁶ In the absence of unobserved heterogeneity, one intuition for why the model is identified is that when there are few potential entrants or entry costs are very low, all potential bidders should choose to enter with high probability. In this case, we have - in essence - exogenous entry and standard results for the identification of value distributions will hold (Athey and Haile, 2006). As potential competition rises or a shifter of entry costs increases, one can then identify the average level of entry costs and the degree of selection from how the amount of observed entry and the value distribution of entrants change.

⁵⁷ It is worth noting that, at the estimated parameters, there is only a single equilibrium because the mills tend to have values that are sufficiently high that only an equilibrium where they are more likely to enter can be supported.

in fixed or sunk entry costs.⁵⁸ This will matter in any setting where the analysis of a policy change, such as liberalizing regulation, divesting plants to a new entrant or allowing a merger, may depend on how strong a competitor a firm that is not currently in the market is likely to be. In Roberts and Sweeting (2012) we consider the effects of selection in the context of airline mergers (making an assumption that corresponds to the FS model laid out above), a setting where it is often argued that the ability of carriers that do not currently serve a route between particular pairs of cities or airports to enter could constrain any route-level market power created by carrier mergers. Of course, this approach could be pushed further to allow for partially selective entry in the way considered here.

One can also raise a broader methodological point about how we should think about entry. One reason why both the theoretical and the empirical auction literatures have been so successful is that the structure of the game that the bidders are playing can be observed and exactly specified within the model. On the other hand, in few settings can we really claim to know the way that the entry game is played, even in a setting like an auction where the ‘market’ that can be entered is clearly defined. When we lack direct information on how a game is played, it is tempting to model it in a way that is as convenient as possible in terms of deriving results, even if the resulting formulation has features that seem unlikely to be true – such as no selection. We hope that the results in this article provide an illustration of how this type of simplifying assumption could lead a researcher to quite misleading conclusions.

References

- Athey, S., Haile, P., 2006. Empirical Models of Auctions. *Advances in Economics and Econometrics. Theory and Applications: Ninth World Congress vol. 2*. Cambridge University Press.
- Athey, S., Levin, J., Seira, E., 2011. Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions. *Q. J. Econ.* 126 (1), 207–257.
- Athey, S., Coey, D., Levin, J., 2013. Set-Asides and Subsidies in Auctions. *Am. Econ. J. Microecon.* 5 (1), 1–27.
- Bajari, P., Hong, H., Ryan, S., 2010. Identification and Estimation of a Discrete Game of Complete Information. *Econometrica* 78 (5), 1529–1568.
- Bhattacharya, V., Roberts, J., Sweeting, A., 2014. Regulating Bidder Participation in Auctions. *RAND J. Econ.* 45 (4), 675–704.
- Bulow, J., Klemperer, P., 1996. Auctions Versus Negotiations. *Am. Econ. Rev.* 86 (1), 180–194.
- Bulow, J., Klemperer, P., 2009. Why Do Sellers (Usually) Prefer Auctions? *Am. Econ. Rev.* 99 (4), 1544–1575.
- Cho, I., Kreps, D., 1987. Signaling Games and Stable Equilibria. *Q. J. Econ.* 102 (2), 179–222.
- Coey, D., Larsen, B., Sweeney, K., 2014. The Bidder Exclusion Effect. Discussion paper. National Bureau of Economic Research.
- Congressional Budget Office, 2005. Small Bidders in License Auctions for Personal Wireless Communications. Discussion paper (Washington DC).
- Cremer, J., Spiegel, Y., Zheng, C., 2009. Symposium on Auction Theory: Auctions with Costly Information Acquisition. *Economic Theory* 38 (1), 41–72.
- Denton, J., 2008. Stacked Deck: Go-Shops and Auction Theory. *Stanf. Law Rev.* 60 (5), 1529–1553.
- Eizenberg, A., 2014. Upstream Innovation and Product Variety in the US Home PC Market. *Rev. Econ. Stud.* 81 (3), 1003–1045.
- Gentry, M., Li, T., 2012. Auctions with Selective Entry. Working Paper.
- Gentry, M., Li, T., 2014. Identification in Auctions with Selective Entry”. *Econometrica* 82 (1), 315–344.
- Gentzkow, M., Shapiro, J., 2014. Measuring the Sensitivity of Parameter Estimates to Sample Statistics. Working Paper.
- Guerre, E., Perrigne, I., Vuong, Q., 2000. Optimal Nonparametric Estimation of First Price Auctions. *Econometrica* 68, 525–574.
- Haile, P., Tamer, E., 2003. Inference with an Incomplete Model of English Auctions. *J. Polit. Econ.* 111 (1), 1–51.
- Halpin, D., 2005. *Construction Management*. 3rd edition. Wiley.
- Hubbard, T., Paarsch, H., 2009. Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation. *Int. J. Ind. Organ.* 27, 1–14.
- Klemperer, P., 2004. *Auctions: Theory and Practice*. Princeton University Press.
- Krasnokutskaya, E., Seim, K., 2011. Bid Preference Programs and Participation in Highway Procurement Auctions. *Am. Econ. Rev.* 101 (6), 2653–2686.
- Krishna, V., 2002. *Auction Theory*. Academic Press, San Diego, CA.
- Levin, D., Smith, J., 1994. Equilibrium in Auctions with Entry. *Am. Econ. Rev.* 84, 585–599.
- Li, T., Zhang, B., 2015. Affiliation and Entry in First-Price Auctions with Heterogeneous Bidders: An Analysis of Merger Effects”. *Am. Econ. J. Microecon.* (forthcoming).
- Li, T., Zheng, X., 2009. Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions. *Rev. Econ. Stud.* 76, 1397–1429.
- Li, T., Zheng, X., 2011. Information Acquisition and/or Bid Preparation: A Structural Analysis of Entry and Bidding in Timber Sale Auctions”. *J. Econ.* 168 (1), 29–46.
- Lu, J., Ye, L., 2013. Efficient and Optimal Mechanisms with Private Information Acquisition Costs. *J. Econ. Theory* 148 (1), 393–408.
- Lu, J., Ye, L., 2014. Optimal Two-stage Auctions with Costly Information Acquisition. Discussion paper. Ohio State University.
- Mankiw, N., Whinston, M., 1986. Free Entry and Social Inefficiency. *RAND J. Econ.* 17 (1), 48–58.
- Marmer, V., Shneyerov, A., Xu, P., 2013. What Model for Entry in First-Price Auctions? A Nonparametric Approach. *J. Econ.* 176 (1), 46–58.
- McAfee, R.P., McMillan, J., 1987. Auctions with Entry. *Econ. Lett.* 23, 343–347.
- Menezes, F., Monteiro, P., 2000. Auctions with Endogenous Participation. *Rev. Econ. Des.* 5, 71–89.
- Milgrom, P., 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- Moreno, D., Wooders, J., 2011. Auctions with Heterogeneous Entry Costs. *RAND J. Econ.* 42 (2), 313–336.
- Myerson, R., 1981. Optimal Auction Design. *Math. Oper. Res.* 6 (1), 58–73.
- Park, W., Chapin, W.B., 1992. *Construction Bidding: Strategic Pricing for Profit*. Second edition. Wiley.
- Quint, D., Hendricks, K., 2013. Indicative Bidding in Auctions with Costly Entry. Working Paper.
- Ramey, G., 1996. D1 Signaling Equilibria with Multiple Signals and a Continuum of Types. *J. Econ. Theory* 69, 508–531.
- Roberts, J., Sweeting, A., 2012. Airline Mergers and the Potential Entry Defense. Working Paper.
- Roberts, J., Sweeting, A., 2013. When Should Sellers Use Auctions? *Am. Econ. Rev.* 103 (5), 1830–1861.
- Roberts, J., Sweeting, A., 2015. Bailouts and the Preservation of Competition”. Working Paper.
- Samuelson, W., 1985. Competitive Bidding with Entry Costs. *Econ. Lett.* 17, 53–57.
- Su, C., Judd, K., 2012. Constrained Optimization Approaches to Estimation of Structural Models. *Econometrica* 80 (5), 2213–2230.
- Tamer, E., 2003. Incomplete simultaneous discrete response model with multiple equilibria. *Rev. Econ. Stud.* 70 (1), 147–165.
- Ye, L., 2007. Indicative Bidding and a Theory of Two-Stage Auctions. *Games Econ. Behav.* 58, 181–207.

⁵⁸ It is common to allow for potential entrants to have additively separable shocks to their payoffs from entering that are not observed by the econometrician. However, these shocks do not affect the payoffs of other players, so they should be understood as shocks to fixed or entry costs, rather than as reflections of the ‘competitiveness’ of the firm that would determine the effect that the entrant has on incumbent profits. In contrast, in our auction model, selection can directly affect whether a new entrant is likely to be the firm that is allocated the object. See Eizenberg (2014) for a recent example of an article assuming that product qualities are only drawn once entry decisions have been taken.