

Set-Asides and Subsidies in Auctions*

Susan Athey, Dominic Coey and Jonathan Levin[†]

October 2010

Abstract

Set-asides and subsidies are used extensively in government procurement and natural resource sales. We analyze these policies in an empirical model of Forest Service timber auctions. The model fits the data well both in-sample and out-of-sample. Our estimates suggest that restricting entry to small businesses substantially reduces efficiency and revenue, although it does increase small business participation. A bidding subsidy for small business appears to be more effective at achieving distributional goals, increasing revenue and eliminating almost all of the efficiency loss. A small business bidding subsidy also increases the average profits of both large and small firms, compared to a set-aside policy. We explain these findings by connecting to the theory of optimal auction design. *JEL Codes:* D44, H57, L53.

*We thank Enrique Seira for his collaboration in assembling the data for this paper, Matt Osborne and Frank Zhuo for research assistance, and Ken Hendricks for comments on an early draft of this paper. The National Science Foundation provided generous financial support.

[†]Athey: Department of Economics, Harvard University, and NBER, athey@fas.harvard.edu; Coey: Department of Economics, Stanford University, dcoey@stanford.edu; Levin: Department of Economics, Stanford University, and NBER, jdlevin@stanford.edu.

1. Introduction

Government procurement programs often seek to achieve distributional goals in addition to other objectives. In the United States, the federal government explicitly aims to award at least 23% of its roughly \$400 billion in annual contracts to small businesses, with lower targets for businesses owned by women, disabled veterans and the economically disadvantaged.¹ Many state and local governments also set goals regarding small businesses or locally owned firms. Given the large scope of these programs, it is perhaps surprising that relatively little is known about the optimal design of preference programs and their costs.

Two common methods are employed to achieve distributional goals. One approach is to *set aside* a fraction of contracts for targeted firms. For instance, federal procurement contracts between \$25,000 and \$100,000 are typically reserved for small businesses.² An alternative is to provide *bid subsidies* for favored firms. Subsidies are used by the Federal government to assist domestic firms bidding for construction contracts under the Buy America Act, by the Federal Communications Commission to favor minority-owned firms in spectrum auctions, and in California state highway procurement to assist small businesses.

This paper develops and estimates a structural economic model of entry and bidding in auctions, and uses it to simulate the revenue and efficiency consequences of using alternative market designs to achieve distributional objectives. Our empirical application is the U.S. Forest Service timber sale program, which conducts both set-aside sales and unrestricted sales, but does not use subsidies. During the time period we study, the Forest Service sold around a billion dollars of timber a year, and in the region from which our data is drawn, 14% of the sales are small business set-asides. We find that designating a sale as a set-aside reduced efficiency by 22% and cost the Forest Service about 16% in revenue. Both losses could have been eliminated, while allocating the same volume of timber to small bidders and increasing aggregate small and large firm profits, by instead providing a 7.5% bid subsidy for small bidders in all sales. If other U.S. procurement and resource allocation programs

¹Section 15(g)(1) of the Small Business Act reads: “The Government wide goal for participation by small business concerns shall be established at not less than 23 percent of the total value of all prime contract awards for each fiscal year.” Extensive documentation of US government procurement programs for small businesses can be found on the Small Business Administration website at <http://www.sba.gov/>.

²See 15 USC 644(g)(1) or the Federal Acquisitions Regulations, Section 19.502-2.

are similar, these results suggest that billions of dollars might be at stake in undertaking a redesign of set-asides.

Basic supply and demand suggests that set-aside programs should lower revenue and decrease efficiency by reducing the number of eligible buyers. This need not be the case, however, if bidding is costly and firms are heterogeneous. In such a setting, restricting participation may increase auction revenue. Suppose there is a single large bidder with a value uniformly distributed between 0 and 30 and two small firms with values uniformly distributed between 0 and 10. If it costs seventy-five cents to learn one's value and enter the auction, the large bidder will be the only entrant and will win at a zero price. If participation is restricted to the small firms, both will enter and the expected price increases from 0 to $3\frac{1}{3}$, despite the fact that expected social surplus decreases by $9\frac{1}{12}$. If there are two large firms, however, both will enter in equilibrium, giving an expected price of 10. In this case, the participation restriction decreases revenue and social surplus. If there are additional large firms, however, or if entry costs are substantially lower or higher, a set-aside program both lowers revenue and decreases efficiency.³ Thus, both entry and bidding behavior must be considered in a full analysis of these programs.

Bid subsidies also can have ambiguous effects depending on the relative strengths of the bidders and the costs of participation. A well-known insight of Myerson (1981) is that appropriately handicapping bidders can increase revenue relative to a standard open or sealed bid auction. The impact of a fixed subsidy, however, can depend subtly on bidders' value distributions, as discussed by McAfee and McMillan (1989). Moreover, with endogenous participation, a subsidy will affect entry in ways that in principle can be helpful or harmful. In the previous example, a rule that awards the object to a small bidder if its bid is at least a third that of the large bidder generates entry by all three firms. Expected revenue increases from zero to $8\frac{1}{3}$ with social surplus decreasing by 4. Such a program results in a small firm winning two-thirds of the time. Less dramatic subsidies have a similar qualitative effect, raising revenue while decreasing surplus. A larger subsidy, however, may discourage participation by the large firm; for some entry costs, the result can be lower revenue than a

³There are other cases where restricting participation can raise revenue even if bidding is not costly, for instance if there are strong "winner's curse" effects; see Bulow and Klemperer (2001).

no-subsidy sale.

Getting a handle on the effect of set-asides or subsidies in a given setting requires understanding the relative strengths of targeted and non-targeted bidders. Forest Service timber sales are characterized by a high degree of diversity in participating bidders. Bidders range from small logging outfits to large vertically integrated forest products companies. We distinguish between the smaller firms that are eligible for set aside sales and the larger firms that are not. The smaller firms are mainly logging companies, while the large firms are mills and often part of larger forest product companies. The relative strength of these bidders varies with the size of the sale. For the smallest quintile of sales by volume, the different types of bidders do not submit significantly different bids. In larger sales, the small firms bid substantially less, and our estimates imply an even greater difference in underlying valuations. One explanation for the bidding and value differences is that large mills can process large quantities of timber more efficiently or avoid frictions in re-selling harvested logs.

To assess the efficiency cost of the set-aside program, we build a model of bidder entry and bidding and estimate its parameters from the data. Building on Athey, Levin and Seira (2011, hereafter ALS), we model each sale as a private value auction with endogenous entry. As we wish to use the model to assess counterfactual changes in the preference policy (unrestricted, set-aside, or subsidy), it is important that it can accurately predict entry and prices when the preference policies change. Thus, we estimate the model using only data from unrestricted sealed bid sales, and assess its performance by comparing the out-of-sample predictions for small business set-asides with the actual outcomes in the data. The model performs well: predicted prices and entry are within 4% of observed values.

Our estimates show that entry responses help to mitigate the losses from set-aside policies. If small bidders did not increase their participation relative to unrestricted sales, revenue and efficiency losses both would be larger (27 and 28 percent, respectively, rather than 16 and 22 percent).

We also use the model to calculate the counterfactual effect of implementing a bidder subsidy program (applied to all sales) in lieu of direct set-asides for a subset of sales.⁴ A range

⁴The idea that the Forest Service set-aside program could be replaced with a subsidy policy is discussed by Froeb and McAfee (1988), and also by Brannman and Froeb (2000).

of subsidies yield outcomes that dominate the observed policy of set-asides. For example, a 7.5% subsidy for small businesses would result in small firms winning at least as much timber as under the set-aside program, but with 4.7% greater program revenue and a 2.7% increase in overall program efficiency. In addition, the average profits of both big and small bidders are increased. So at least on our setting, modestly subsidizing weaker bidders in all auctions appears preferable to excluding strong bidders from some auctions. This attraction of subsidies can be understood by connecting our empirical model to the theory of optimal auction design, which we do in the final section of the paper.

Another way to achieve distributional objectives with smaller efficiency losses is to select sales to be set-asides where efficiency losses would be small. We build a simple empirical model that selects sales into the set-aside program in a way that minimizes expected efficiency losses subject to a constraint of volume sold, and find that using this model to allocate sales into the set-aside program would result in revenue and efficiency that are virtually identical to the no-preference policy.

We also investigate the idea that a set-aside program serves to “guarantee” a minimal level of timber for targeted firms, reducing the risk that small bidders will win little timber. However, we find that this benefit is modest due to the relatively large number of sales.

Our results can be usefully compared to recent findings of Marion (2007) and Krasnokutskaya and Seim (2010), who study the effect of bid subsidies in California highway procurement auctions.⁵ Marion compares state-funded auctions that have a small businesses subsidy to federally-funded auctions with no subsidy. He finds that procurement costs are 3.8% higher in the subsidy auctions, and attributes the increase to decreased participation by large firms in subsidy auctions. Krasnokutskaya and Seim use data from the subsidy auctions to estimate a structural bidding model, and use the model to simulate alternative preference policies. They conclude that the subsidy program has a very small effect on procurement costs, less than 1%.

⁵Several other papers also simulate various types of preference policies as applications of estimated auction models. Examples include Brannman and Froeb (2000), Flambar and Perrigne (2008), Roberts and Sweeting (2010). Brannman and Froeb’s paper, which looks at Forest Service timber auctions, is particularly interesting because although the approach is quite different from ours (they do not consider bidder participation, use different data, and consider a logit value model of second price auctions), they reach a similar conclusion about the revenue effect of the Forest Service set-aside program.

An intermediate finding in these papers, and one that contrasts with our setting, is that the large firms in the California highway auctions do not appear to have much of a cost advantage, so the “Myerson effect” of subsidies is small. Another difference is that we estimate a complete model of entry and bidding using data on non-set-aside auctions, and establish that our model provides accurate predictions (out of sample) of the outcomes in small business set aside sales, providing greater confidence in our counterfactual simulations. That being said, all three studies share a central theme, which is that accurately accounting for participation is crucial in assessing bid preference programs.

2. A Model of Set-Asides and Subsidies

This section describes our basic model of the auction process, which builds on ALS. We then use the model to informally discuss the effect of set-asides or bidder subsidy programs.

A. The Model

Consider a seller who wishes to auction a single tract of timber. She announces a reserve price r , and whether the auction will be open or sealed bid. There are N_S potential small bidders and N_B potential big bidders. The potential bidders have values that are independently distributed according to either F_S or F_B depending on the bidder’s size. These distributions have densities f_τ and supports $[0, \bar{v}_\tau]$ for $\tau = S, B$. A bidder must spend K to learn its private value and enter the auction. After the entry decisions are made, each participant learns the identities of the other participants before bids are submitted. In a sealed bid auction, the highest bidder wins and pays its bid. In an open auction, the highest bidder wins and pays the second highest bid (or the reserve price if there is a single bidder).

The analysis of the bidding game is standard, so we focus on the entry game. Let $\pi_S(n_S, n_B)$ and $\pi_B(n_S, n_B)$ denote the expected profit of an entering small or large bidder, gross of the entry cost, as a function of the number of small and big participants (n_S, n_B) . In a pure strategy entry equilibrium, each participant has an expected profit of at least K , and if there are non-participating bidders, the expected profit to these bidders of entering is less than K . There are two cases of particular interest. In the first case, which will correspond to the smallest sales in our data, the small and big bidders have identical value distributions,

so $F_S = F_B$. In this case, there can be multiple entry equilibria, but the total number of participants is the same in all of them. In the second case, which will describe the majority of sales in our data, the big bidders have a higher value distribution. For this case, ALS point out that if the big bidders are sufficiently strong relative to the small bidders, then for all n_S, n_B , we will have $\pi_B(n_S, n_B + 1) > \pi_S(n_S, n_B)$. When this “strong asymmetry” condition holds, the entry equilibria have a convenient form. Either no small bidders enter in equilibrium, or if small bidders do enter, then all of the potential big bidders also enter.

Both in our empirical estimation and in simulating auction outcomes, we will focus on *pure strategy* entry equilibria. The model may also have mixed strategy equilibria, which have some appeal because “type-symmetric” equilibria (i.e. equilibria with common strategies for all big bidders, and for all small bidders) are typically in mixed strategies. ALS focus on type-symmetric equilibria in their comparison of open and sealed bidding, although they point out that analyzing pure strategy equilibria would be similar. The main reason we focus on pure strategy equilibria here is that if one considers mixed equilibria, a decrease in the number of potential bidders — for example, due to a set-aside program — can increase expected participation and revenue even if bidders are symmetric.⁶ In our particular setting, we view this “coordination effect” as relatively implausible, so focusing on pure strategy equilibria seems like a reasonable choice.

We should be clear, however, that modeling entry requires a number of modeling decisions that can be debated. In addition to the issue of whether to focus on pure or mixed equilibria, or symmetric equilibria, or all possible equilibria, one can ask if bidders have private information about their values prior to making entry decisions, whether they have or acquire common information that is unobserved to the econometrician, and whether they can acquire information about the level of likely competition. Recent papers that take different approaches to these issues include Li (2005), Li and Zheng (2009), Bajari, Hong and Ryan (2009), Krasnokutskaya and Seim (2010), Marmer, Shneyerov and Xu (2010) and Roberts and Sweeting (2010). Several of these papers highlight the coordination benefit of restricting

⁶Consider an auction with three potential bidders who have values distributed uniformly on $[0,10]$ and an entry cost of $10/3$. With no entry restriction, all bidders will enter with probability $2/3$ in equilibrium and expected revenue will be $80/27$. If one potential bidder is restricted from entering, the two remaining firms will enter with probability 1 and expected revenue will increase $10/3$.

potential entry in a mixed strategy equilibrium, which we de-emphasize. The last two papers allow bidders to have a degree of private value information before making their entry decisions, and use variation in potential entry to aid identification.⁷ On the other hand, they make some assumptions that are not ideal for our purposes: Marmer, Shneyerov and Xu (2010) assume that all bidders are symmetric and rule out unobserved heterogeneity across auctions, while Roberts and Sweeting’s (2010) approach, at least in its current version, is applicable only to open auctions.

B. Set-Aside Auctions

A small business set-aside excludes big bidders from the auction, but increases the incentives for small bidders to participate because they anticipate less competition. If the small and big bidders have identical value distributions, and there are a sufficient number of potential small entrants to substitute one for one for big entrants, a set-aside provision will have no effect on total participation, revenue or overall efficiency. If there are not enough potential small entrants to promote fully compensating entry, a set-aside will reduce total entry, reduce revenue, and reduce efficiency.

The effects of a set-aside are less clear if the value distributions are asymmetric. The reason is that, as explained in the introduction, the increase in small bidder participation can lead to a greater overall number of auction entrants, and potentially, to higher expected revenue. The effect of a set-aside, both in direction and size, therefore depends on the number and relative strength of potential bidders, the cost of entry, and the auction format. For the open auction case, it is possible to show that a set-aside cannot increase total surplus because the unrestricted entry equilibrium is socially efficient — that is, it maximizes expected surplus given the set of potential entrants and simultaneous entry decisions.

C. Bidder Subsidies

A subsidy program favors the bids of certain firms. A typical approach is to say that a favored bidder must pay only a portion $b/(1 + \alpha)$ of its bid b for some $\alpha > 0$. (One can also make the bid credit an absolute amount rather than a fraction of the bid.) To illustrate the

⁷This requires a good estimate of potential entry, which is not the strongest aspect of our data. On the other hand, our data has excellent information on realized entry, even for open auctions, which is something that our approach exploits.

effects of a subsidy, suppose we have an open auction with two participants, a big bidder with value v_B and a small bidder with value v_S . If the seller offers a subsidy of size α to the small bidder, there are three possible outcomes. If $v_S > v_B$, the subsidy will not change the outcome of the sale, but it will lower revenue from v_B to $v_B/(1 + \alpha)$. If $(1 + \alpha)v_S > v_B > v_S$, the subsidy will allow the small bidder to win over the higher-valued big bidder and revenue will fall from v_S to $v_B/(1 + \alpha)$. Finally, if $v_B > (1 + \alpha)v_S$, the big bidder will win with or without the subsidy, but the policy will raise revenue by αv_S .

From an ex-ante standpoint, it is relatively easy to see that if big bidders are stronger a small subsidy will tend to increase sale revenue. A small subsidy is unlikely to affect the allocation and conditional on the allocation being unaffected the big bidder is the likely winner, so revenue increases. A similar logic applies even if there are more bidders, although the subsidy can end up being irrelevant or neutralized if the high bidders are both small or both big. A small subsidy will also increase small bidder participation without affecting the participation of large bidders, which leads to another positive revenue effect, although at the cost of distorting social efficiency.

The situation becomes ambiguous if one considers larger subsidies. For fixed participation the allocative distortions become larger, and strong but unfavored bidders also may be deterred from participating. So in principle, some subsidies may reduce both revenue and social efficiency.

3. Description of Timber Sales

This section describes how timber auctions worked in the time period we consider, the small business set-aside program, and the data for our study. We discuss only the essentials of the sale process; more detailed accounts can be found in Baldwin, Marshall and Richard (1997), Haile (2001), Athey and Levin (2001) or ALS.

A. Timber Sales and Small Business Set-Asides

A sale begins with the Forest Service identifying a tract of timber to be sold and conducting a survey to estimate the quantity and value of the timber and the likely costs of harvesting. A sale announcement that includes these estimates is made at least thirty days

prior to the auction. The bidders then have the opportunity to conduct their own surveys and prepare bids. The Forest Service uses both open and sealed bid auctions. If the auction is open, bidders first submit qualifying bids, typically at the reserve price, followed by an ascending auction. The sealed bid auctions are first price auctions. In either case, the auction winner has a set period of time, typically between one and four years, to harvest the timber.

The Forest Service designates certain sales as small business set-asides. For a standard set-aside sale, eligible firms must meet two basic criteria. First, they must have no more than 500 employees. Second, they must manufacture the timber themselves or re-sell it to another small business, with the exception of a specified fraction of the timber for which no restrictions apply. In our data, there appear to be some exceptions to the eligibility criteria, and conversations with Forest Service employees confirm that the rules are occasionally loosened for various reasons.

The Forest Service regulations also provide guidelines for which sales should be designated as set-asides.⁸ The Forest Service periodically sets targets for the amount of timber small businesses are expected to purchase in different areas. Though subject to some adjustment, the basic goal is to maintain the historical share of timber volume logged by small businesses in different areas, with the historical amounts corresponding to the quantities logged between 1966 and 1970. By projecting the amount of timber that will be purchased by small businesses in unrestricted sales, the Forest Service determines the quantity of timber that must be sold using set-aside sales, although forest managers have some discretion to accommodate specific local needs. Forest managers are expected to use the same sale methods for set-aside sales and to include a variety of sale sizes, terms and qualities in the set-aside program. Forest managers do have some discretion to designate tracts as set-asides based on the needs of small businesses in the area, which raises the possibility that tracts designated as set-asides may be relatively well-suited to small firms. We revisit this below.

B. Data and Descriptive Analysis

Our data consists of sales held in California between 1982 and 1989. For each sale, we know the identity and bid of each participating bidder, as well as detailed sale characteristics

⁸See the U.S. Forest Service Handbook, Section 2409.18 on Timber Sale Preparation.

from the sale announcement. We also collected additional information to capture market conditions. We use national housing starts in the six months prior to a sale to proxy for demand conditions, and U.S. Census counts of the number of logging companies and sawmills in the county of each sale as a measure of local industry activity. Finally, for each sale, we construct a measure of active bidders in the area by counting the number of distinct firms that bid in the same forest district over the prior year.

We use participation in set-aside sales, combined with internet searches on individual firms, to construct an indicator of small business status for each firm. ALS distinguishes between mills that have manufacturing capability and logging companies that do not. Essentially all of the logging companies are small businesses. The largest and most active mills are not, but there are also some smaller mills that are eligible for set-asides.⁹ In this paper, we classify bidders based on their small-business status (and refer to them as small and big) rather than their manufacturing capability. An alternative would have been to treat small mills as a separate category, but in later simulations that require us to compute sealed bid auction equilibria, the inclusion of a third type of bidder adds additional complication to the already challenging problem of accurate computation. We provide some evidence additional discussion of the small business classification in Appendix A.

Table 1 presents summary statistics of tract characteristics and auction outcomes for unrestricted and set-aside sales. The participation variables suggest that although set-asides decrease the number of eligible bidders, this does not translate directly into a fall in realized total participation. Additional participation by logging companies substitutes for the absent mills. The table also shows that for both big and small tracts, prices are somewhat lower when participation is restricted, although the difference is not statistically significant.

Data from unrestricted sealed sales suggests that the bidding behavior of big and small bidders is on average quite different. To examine this more closely, we first stratify sales by the volume of timber being auctioned. We then regress the logarithm of per-unit sealed bids on auction fixed effects and an indicator variable for big firms. We do this separately for sealed bid auctions in each sale size quintile. Table 2 shows the results. For the smallest

⁹We observe a few set-aside sales in which large mills entered, presumably because of exceptions made to the rules. There are nine of these sales and we drop them from the analysis.

sales, the estimated difference between the bids of small and large firms is not statistically or economically significant, while big firms bid about 11% more in second quintile tracts. The coefficients for larger tracts are imprecisely estimated because these sales are predominantly open auctions, but the final column of the table shows that if we consider both the open and sealed sales, larger sales are more likely to be won by big bidders. In our empirical model below, we therefore allow the asymmetry between big and small bidders to depend on the size of the sale, with no asymmetry for sales in the smallest quintile by volume, and asymmetry for larger sales.

A key issue for our empirical analysis is the extent to which the tracts designated as set-asides differ from those where participation is unrestricted. The tracts should be comparable within a given forest based on Forest Service regulations, but forest managers also have some discretion. Table 1 also indicates that at least on observable characteristics there are not large differences. To explore the point further, we use a logistic regression to estimate the probability that a sale is set aside as a function of observable tract characteristics. The results appear in Table 3. The most economically and statistically significant explanatory variables are the forest dummies, indicating that the use of set-asides varies across forest, consistent with the USFS policy to preserve historical volumes allocated to small bidders. Sales with higher logging costs (perhaps requiring more complex equipment) are less likely to be small business set-asides. We control for these tract characteristics in our empirical models.

We then consider whether forest managers might designate as set-asides tracts that are relatively more attractive to small bidders. We use the logit estimates to compute the estimated probability that each tract is designated a set-aside; we refer to this as the “set-aside propensity score.” We then consider the tracts that were *not* designated as set-asides and estimate a logit regression to estimate the probability that the sale is won by a small bidder, including the set-aside propensity score along with other sale characteristics as an explanatory variable. The propensity score is not significantly related to the type of bidder that wins the auction in either an economic or statistically significant way (Appendix Table A5), providing some evidence that set-asides are not designated on the basis of their attractiveness to small bidders.

As a first pass at assessing the effect of set-asides, we consider the linear model:

$$Y = \delta \cdot SBA + X\beta + SBA \cdot X\gamma + \varepsilon. \quad (1)$$

Here Y is an outcome of interest (total participation or $\log(\text{revenue})$), SBA is a dummy equal to one if the sale is a small-business set-aside, and X is a vector of observed sale and forest characteristics, including the propensity score from the logit regression described above. We expect the set-aside effect to vary across auctions, particularly as a function of sale size, so we allow for alternative interaction effects in our specifications. The key assumption for identification is that the choice of whether to make the sale a set-aside is uncorrelated with unobservables that might directly affect the outcome, that is ε and SBA are independent conditional on X .

We report regression results in Table 4. The reported estimates represent the average effect of set-aside status across tracts in the sample, or $\hat{\delta} + \bar{X}\hat{\gamma}$, where \bar{X} is the vector sale covariate mean values). They estimates indicate that set-aside status is associated with a fall in revenue, and a reduction in overall entry, although the standard errors are relatively large reflecting the modest sample size.

The structural model we estimate below allows us to go beyond these preliminary regressions in two ways. First, we incorporate information on losing bids into the estimation, which narrows the uncertainty in the estimated effects. Second, it gives us a framework for identifying the channels through which restricting entry affects outcomes, as well as for analyzing welfare and for evaluating a range of counterfactual subsidy policies.

4. Estimating Economic Primitives

We now calibrate our theoretical model from Section 2 and use it to analyze the impact of different policies. The primitives of the model are the value distributions of the big and small bidders, the entry cost for each auction, and the numbers of potential entrants. We first estimate the value distributions (as a function of sale characteristics) from the bid distributions in unrestricted sealed auctions. We then estimate entry costs and potential entrants for each auction as a function of sale characteristics and entry patterns in the local

geographic area. In Section 5, we use the estimated model to investigate the effect of the set-aside program and to study the potential impact of bidder subsidies.

A. Bidders' Value Distributions

Our approach to estimating the bidders' value distributions follows ALS. They employ a version of the two-step methods developed by Guerre, Perrigne and Vuong (2000) and Krasnokutskaya (2009). The first step fits a parametric model of the bid distributions in sealed auctions, allowing these distributions to depend on observed sale characteristics and on an unobserved sale characteristic that accounts for within-auction correlation of bids.¹⁰ The second step uses nonparametric methods to estimate the implied value distributions. Estimates of these primitives allow us to compute expected bidder profits conditional on entry, and consequently to infer entry costs.

Let $F_\tau(\cdot|X, u)$ denote the value distribution for a bidder of size $\tau \in \{S, B\}$ conditional on the observed sale characteristics X , and an unobserved sale characteristic u . We assume that (X, u) and the number of actual participants $n = (n_S, n_B)$ are common knowledge to the bidders at the time they submit their bids. Consistent with our theoretical model, we assume that bidder values are independent conditional on (X, u) and that participants use equilibrium bidding strategies. If there is a single entrant to the auction, we assume he bids the reserve price. With multiple entrants, we write the equilibrium bid distributions as $G_\tau(\cdot|X, u, n)$.

Based on our earlier observation that bidder heterogeneity appears to matter for most sales, but not for the very smallest sales, we split the sample by sale size to estimate the equilibrium bid distributions. That is, we distinguish the small tracts with timber volumes in the lowest quintile from the larger tracts.

In each case, we follow ALS and assume that the unobserved auction characteristic u is drawn from a Gamma distribution with mean one and variance θ , independent of X and

¹⁰Krasnokutskaya was the first to point out the importance of allowing for unobserved auction heterogeneity in estimating auction models, and how the Guerre et al. approach for sealed bid auctions could be extended in this direction. In principle, the open auction data also contain relevant information about bidder values, but the open outcry nature of these auctions is a complicating factor. See ALS for a discussion that summarizes points made by Athey and Haile (2002) and Haile and Tamer (2003). Roberts and Sweeting (2010) use open auction data to estimate a model with unobserved heterogeneity by making use of parametric assumptions.

n . We assume that conditional on (X, u, n) , the bids of small and big firms have a Weibull distribution.¹¹ That is, for $\tau = S, B$,

$$G_\tau(b|X, u, n) = 1 - \exp\left(-u \cdot \left(\frac{b}{\lambda_\tau(X, n)}\right)^{\rho_\tau(n)}\right). \quad (2)$$

In equation (2), $\lambda_\tau(\cdot)$ is the scale of the Weibull distribution, parameterized as $\ln \lambda_\tau(X, N, n) = X\beta_X + n\beta_{n,\tau} + \beta_{0,\tau}$, while $\rho_\tau(\cdot)$ is the shape, parametrized as $\ln \rho_\tau(n) = n\gamma_{n,k} + \gamma_{0,k}$. For small sales, the bid distributions of the small and big firms are modelled as symmetric.

Table 6 reports estimated coefficients of the bid distribution parameters (the summary statistics for the estimation sample are reported in Appendix Table A2). There is strong evidence for unobserved auction heterogeneity, indicated by the estimated variance parameter θ . In the larger sales, the bids of the big firms stochastically dominate those of the small firms. Bids are also increasing in the number of competitors.

Given estimates of the equilibrium bid distributions, we follow Guerre, Perrigne and Vuong, and Krasnokutskaya, in inferring the bidders' value distributions. If the bids in the data are generated by equilibrium bidding, then in an auction with characteristics (X, u, n) a bidder i 's bid b_i and his value v_i are related by the first order condition for optimal bidding:

$$v_i = \phi_i(b_i; X, u, n) = b_i + \frac{1}{\sum_{j \in n \setminus i} \frac{g_j(b_i|X, u, n)}{G_j(b_i|X, u, n)}}. \quad (3)$$

Having estimated each G_j , the only difficulty in inferring values is that we do not observe the unobserved sale characteristic u corresponding to each observed bid. We do, however, have an estimate of its distribution, so we can infer the distributions $F_S(\cdot|X, u)$ and $F_B(\cdot|X, u)$ for any value of u from the relationship:

$$F_\tau(v|X, u) = G_\tau(\phi_\tau^{-1}(v; X, u, n)|X, u, n).$$

¹¹Athey, Levin and Seira discuss the motivation both for using a parametric model of the bid distributions and for the specific choice of the Gamma-Weibull functional form. It is possible to test the appropriateness of the parametric assumption using Andrews' (1997) Conditional Kolmogorov test. Using this test, we cannot reject the null hypothesis that the parametric model accurately describes the data at a 27% level for big sales, and at a 31% level for small sales.

Two subtleties arise in this step. First, if bidders' value distributions are to be bounded, the equilibrium bid distributions must be as well, in contradiction to our Weibull specification. We follow ALS's procedure and truncate the estimated bid distributions. Second, the theoretical value distributions do not depend on the actual number of bidders n , but as is typical in two-stage estimation of auction models, there is some variation in our estimated distributions. There are several approaches to this problem. One is to average the estimated value distributions obtained for different value of n . The problem is that this leads to an average bid distribution that is more spread out than any one distribution, which in turn leads us to infer unrealistically high markups. Instead, we use the estimated value distribution corresponding to $n_S = n_B = 2$, this being a particularly common entry combination. To provide a rough sense of the relationship between bids and values, we calculate that with two small entrants and two big entrants, the median sealed bid markup varies from 12.6% to 14.5% depending on the size of the sale and the type of bidder. These figures are comparable to those reported in ALS, who used a somewhat different dataset that did not include set-aside sales, but included salvage sales.

B. Bidder Profits as a Function of Entry and Sale Characteristics

Given the estimated value distributions, we can find bidder profits as a function of (X, u, n) by simulating either open or sealed bid auctions. The simulation procedure works as follows. We first use the estimated value distributions, $F_S(\cdot|X, u, n)$ and $F_B(\cdot|X, u, n)$ to compute the expected profits of a small and big entrant as a function of (X, u, n) . This step is straightforward for an open auction because the equilibrium strategy is simply to bid one's value, so expected outcomes can be calculated by repeatedly drawing bidder values and calculating auction outcomes. For sealed bid auctions, the simulation is similarly straightforward because we have already estimated the inverse equilibrium bid functions ϕ_S, ϕ_B above (as described in equation (3)).

Given estimates of the profit functions $\pi_S(X, u, n)$ and $\pi_B(X, u, n)$ we average over values of u , according the estimated distribution G_U . This gives us the expected small and big bidder profits $\pi_S(X, n)$ and $\pi_B(X, n)$ that are relevant for the entry decision. (The expected profits also depend on whether a sale is open or sealed bid, but this will not be

made explicit unless necessary.) For the small sales, the expected profits of small and big entrants are the same, due to the symmetric value distributions. For the larger sales, big bidders have higher expected profit. For example, in a large open auction with two small and two large bidders, the expected profit of a big is 3.36 times the expected profit of a small bidder; the ratio is 3.23 in a sealed auction. In fact, for all but 38 of the larger sales in the data, expected profits satisfy the strong asymmetry condition defined earlier. As noted above, this condition is useful for narrowing the range of possible entry equilibria.

C. Entry Cost and Potential Entrants

We next use our estimates of bidder expected profits to infer entry costs from the observed entry pattern. Assuming that the observed entry behavior corresponds to a pure strategy entry equilibrium, the estimated bidder profits imply both upper and lower bounds on the entry cost for each sale. To see how this works, consider what is revealed by the decision of bidders to enter a sale with characteristics X and observed (non-zero) entry $n = (n_S, n_B)$. If a small bidder entered, so $n_S > 0$, we can infer that $K \leq \pi_S(X, n)$. If not, we can still infer that $K \leq \pi_B(X, n)$. These inequalities provide an upper bound on the entry cost of each sale. Additional information can be obtained from the decisions of potential entrants not to participate. If we observe a small bidder not entering, we can infer that $K \geq \pi_S(X, n_S + 1, n_B)$. Similarly, if we observe a big bidder not entering, then $K \geq \pi_B(X, n_S, n_B + 1)$. The small bidder inequality provides the more conservative (i.e. lower) lower bound.

Implementing this procedure for each sale leads to two challenges. First, the expected profits of entrants (the upper bound on entry costs) are often too high to convey much meaningful information. For instance, five dollars per thousand board feet (\$5/mbf) would appear to be a fairly cautious upper bound on the cost of surveying the tract. Yet the estimated expected profits of the observed entrants exceed this level in 50% percent of the sales. For the lower bound, determined by the decisions of the non-entrants, we face a different challenge. We require an accurate measure of potential entry. To see why, consider what happens if we observe an auction for which $\pi_S(X, n_S + 1, n_B)$ is large compared to any plausible entry cost. While the most natural way to rationalize this might be to conclude

that the set of potential small entrants was exhausted, relying on an imperfect measure of potential entry can lead to the conclusion that entry costs were extraordinarily high for that particular sale.¹²

We address this problem by bringing to bear auxiliary information, and combining it with an imperfect measure of the potential small entrants. We start by using the maximum number of small bidders who entered a sale in a given forest district as a measure of potential entry in that district (there are on average 29 sales per forest district in our sample). We then compute the expected profits of the eligible non-entrants for each sale. For 50% of the sales, this number is below \$5/mbf. In these cases, we treat the number of potential entrants as valid and use the estimated profit of the non-entrants as the implied entry cost. For the remaining sales, we conclude that potential entry was exhausted so that we cannot learn about the lower bound on entry costs. Instead we extrapolate a lower bound on entry cost using a regression model estimated on the 50% of sales where the lower bound can be inferred directly. (Appendix Table A3 presents the estimates from this model.) We use the same model to predict entry costs for sales outside the estimation sample.¹³

Because this procedure is a practical but imperfect way to address a data limitation, we investigated its robustness and experimented with other approaches as well. In one check, we examined whether our later simulation results were sensitive to the particular \$5/mbf assumption. We found little sensitivity to using values as low as \$2.50/mbf or as high as \$10/mbf. The choice affects our estimate of potential small bidder entry, but the comparison between set-asides and subsidies is robust to these changes. This is unsurprising: the effect of a fifth or sixth potential small bidder is small relative to the effect of removing all big bidders. We also considered (type-symmetric) mixed strategy entry equilibria, but we have already discussed our concerns with the comparative statics properties they create with regard to set-asides. Finally, we attempted to estimate a parametric joint distribution of potential entrants and entry costs. We viewed this as initially appealing, but found it

¹²Athey, Levin and Seira note that measurement error in potential entrants may be less of a problem if one focuses on mixed strategy entry equilibria, because the key statistical variable is the observed (binomial) distribution of competing entrants, and in fitting this distribution, there is not a high degree of sensitivity to small changes in the number of potential entrants. As noted above, we focus on pure strategy equilibria in this setting because of comparative statics properties of the mixed equilibria.

¹³For three sales, the projected entry cost from the regression exceeds the expected profit of an entrant (the upper bound on entry costs). In these cases, we use the upper bound as the estimated entry cost.

difficult to obtain estimates of entry costs that seemed plausible without imposing auxiliary assumptions similar to the one above. This suggests that the data alone make it difficult to distinguish sales where a potential entrant was deterred by high entry costs from sales in which entry costs were lower but the set of potential entrants was exhausted.

The procedure we have described gives us an estimated upper and lower bound on the entry cost at each sale. In our simulations, we use the lower bound estimates. The main reason to use the lower bounds is that they are still fairly high: \$10,972 for the median sale. The one remaining variable to be inferred is the number of potential big entrants. For the larger unrestricted sales in our data, the strong asymmetry of bidder values means that if a small bidder entered, then if we are observing equilibrium entry, all big bidders must have entered. For these sales, we infer that the number of potential big entrants was equal to the number of actual big entrants. For a second set of sales with unrestricted entry, the estimated entry cost is below the estimated expected profit of an additional big entrant. For these sales, we also conclude that the number of potential big entrants equals the number of actual big entrants. In the 3% of sales which fall into neither of these two categories, we cannot infer the potential big bidder entry directly. We assume that the number of potential big bidders is Poisson distributed conditional on characteristics X , estimate the parameters on the sales where potential big bidder entry can be directly inferred, and then for the other sales take draws from the implied distribution of potential big entrants when computing simulated outcomes. We also use the Poisson model to estimate potential big bidder entry as a function of sale characteristics for sales outside our estimation sample.

5. Analysis of Set-Asides and Subsidies

In this section, we use the estimated model to analyze the set-aside program, and to evaluate whether a subsidy program could achieve the government's distributional goals at lower cost in terms of revenue and efficiency.

A. Assessing the Fit of the Model Inside and Outside the Estimation Sample

We use our model to predict the outcome of a small business set-aside auction and an unrestricted auction for each tract in our data sample. For the subset of tracts where the

actual sale was unrestricted, we compare the model’s prediction for an unrestricted auction to assess the model’s fit in-sample.¹⁴ For the tracts where the actual sale was a set-aside, we compare our prediction of the set-aside outcome as a way to assess the model’s ability to make out-of-sample predictions. (Note that by out-of-sample, we mean both on tract characteristics and entry restrictions — the model was estimated only using the unrestricted sales.) Throughout, we hold fixed the auction format: we simulate sealed bid auctions if a tract was sold by sealed bidding, and open auctions if a tract was sold by open auction.

When we make counterfactual predictions, we need to predict entry given our estimates of profits conditional on entry, costs, and potential entrants. In general, there are multiple equilibria and the number of small and big entrants can vary across these equilibria. For instance, if we observe a small sale (where value distributions are symmetric) with $N_S = N_B = 2$ and there is an equilibrium with $(n_S, n_B) = (2, 0)$, then $(1, 1)$ and $(0, 2)$ also must be equilibrium entry configurations. We accommodate this by computing bounds on the equilibrium outcomes (such as the probability of a small firm winning the auction). The number of small and big entrants vary across equilibria, however, in less than 5% of the sales in our data, so the width of the bounds is negligible — less than 1/5 of one percent for the estimated outcomes. Because of this, we report just the midpoint.

The results are reported in Table 6, which is divided into two panels. The first reports average auction outcomes for the tracts that were sold in unrestricted sales. The second reports the same outcomes for small business set-aside tracts. In each case, we report the sale outcomes observed in the data, and the model’s predicted outcomes from running the sales in unrestricted fashion and as small business set-asides. Comparing the first and second columns in the top panel illustrates the model’s fit in sample: predicted prices are within 3% of the actual prices, and the percentage of sales won by small bidders is within 4% of the actual value.

The more challenging test for the model is the out-of-sample fit for tracts that were sold as set-asides. The results can be seen by comparing the first and third columns of the bottom panel. The model predicts entry within 4% of the actual value, prices within 2% and the

¹⁴Note that the bidder value distributions are estimated only on the sealed bid unrestricted sales, but the entry costs are estimated using both open and sealed bid unrestricted sales.

percentage of sales won by small bidders within 1%.

B. Assessing the Impact of Set-Asides

Table 6 also illustrates the impact of the set-aside policy on revenue, entry and welfare. For the sales that were in fact set-asides, the model suggests that opening up entry would have resulted in around 1.66 big bidders per sale, with one fewer small entrant. The fraction of sales won by small bidders would have dropped to 60%, but revenue would have been 19% higher, and surplus 28% higher. The model yields roughly similar predictions for the effect of a set-aside on the tracts that were sold in unrestricted fashion. A set-aside would have reduced the entry of big firms to zero, and the model suggests that the increase in small bidder entry would not have fully compensated in terms of revenue, leading to a 13% revenue loss. The predicted surplus loss is again greater, at 21%.¹⁵ These results are driven almost entirely by the improved efficiency in allocation. The reduced entry costs from unrestricted sales contribute less than 1% of the gain in surplus.

C. Assessing the Impact of Subsidies

We next use the model to consider an alternative small-business subsidy policy that might substitute for the set-aside program. In particular, we consider a policy under a small businesses would pay only $1/(1 + \alpha)$ of its bid if it won an auction. We ask whether there are values of α for which the Forest Service could increase revenue and economic efficiency while selling the same fraction of timber to small businesses.

To simulate the effect of subsidies, we use essentially the same procedure as described in the previous section. There is, however, an important qualification. Because the equilibrium sealed bidding strategies in a subsidized auction do not correspond to objects that we observe directly in the data, we need to compute them. Computing equilibrium strategies in asymmetric auctions is well-known to be a challenging problem (Marshall et al., 1994; Bajari, 2001). The approach we take is to solve for equilibrium bidding strategies in an auction in which the bid space is discretized and use this to approximate the equilibrium

¹⁵For the tracts sold by unrestricted sale, a simulated set-aside policy sometimes lead to having zero entrants. This occurs for 86 of the 1167 unrestricted sales. The reported price and surplus numbers are conditional on having positive total entry. If we assume that the Forest Service gets zero surplus from the unsold tracts, the predicted surplus per sale falls to 319.89.

in the underlying game with a continuous bid space. The details of the computation are provided in Appendix B, and are perhaps of some independent interest. The validity of the approximation derives from Athey (2001), who shows that for a class of auction games that includes our model, the equilibria of auctions in which the bidding range is divided into a grid converge to an equilibrium of the continuous bid space game as the grid becomes finer.

Table 7 reports expected outcomes, averaged over all auctions in the sample, and then with the smallest sales broken out, for five different subsidy levels (no subsidy or $\alpha = 0$, 2.5%, 5%, 7.5% and 10%). The model's predicted outcomes can be compared to the actual outcomes, which are also reported in the Table. The first point to make is that a relatively low subsidy level, no more than 5%, appears sufficient to ensure that small businesses win the same fraction of sales and volume of timber as under the observed set-aside policy. The second point is that in our simulations, these subsidies increase revenue and efficiency over existing policy. In fact, revenue is increasing in the subsidy level, at least up to a 10% subsidy. The revenue effect arises because the subsidy makes small firms more competitive in the large sales (partially mimicking the allocation rule of an optimal auction), with the additional effect of increasing small business participation. The model predicts that a 10% subsidy would increase average small business participation in the larger sales from 2.78 to 3.27.

A third point that comes out of the simulation results is that there appear to be subsidy levels that result in fairly widespread benefits relative to the set-aside policy. For example, a 7.5% subsidy entails more total surplus than the observed set-aside policy and in fact, 99.9% as much surplus as a no-subsidy, no set-aside case. This subsidy level also results in slightly more sales and timber won by small firms, greater profits for both types of bidders, and greater revenue than the set-aside policy. Even a relatively large 10% subsidy still leads to greater profits for big firms because they are not excluded completely from certain sales. A greater subsidy at this level does begin to reduce efficiency, in part by encouraging excessive entry. Conditional on entry, small subsidies induce small efficiency costs: in an open auction, for example, the allocative inefficiency is bounded above by the level of the subsidy. Overall, however, the distortions created by a subsidy policy appear to be relatively small compared to the costs of excluding high-value big firms from set-aside sales.

Fourth, we examine whether the benefits of subsidies can be approximated by using a better-designed set-aside policy. Since the evidence suggests that small bidders are attracted to small sales, we first consider a program that allocates all of the smallest sales to be set-asides. However, the volume-based approach leads to lower efficiency and revenue than the existing set-aside program. We then consider a more sophisticated alternative. We predict the sales that would be most efficient to designate as set-asides given the constraint that small bidders win as much volume as in the existing program. To do this, we select the sales that have the fitted values closest to zero in a regression of (Unrestricted Surplus - Set-Aside Surplus)/(Unrestricted Small Bidder Volume - Set-Aside Small Bidder Volume) on tract characteristics, selecting just enough sales that the volume constraint is satisfied. This approach leads to outcomes almost as efficient as under a subsidy policy, and revenue is approximately the same as if there were no preference program at all, though 2.4% less than with a 7.5% subsidy. One disadvantage is that small bidder profits are 10% lower than with either the existing policy or the 7.5% subsidy (though large bidder profits are higher).

Table 8 decomposes the impact of subsidies and set-asides into bidding and participation effects. We find that for sales that were sold as set-asides, endogenous entry reduced the loss in surplus (over an unrestricted policy) from 28% to 22%, and it reduced the revenue losses from 27% to 16%. For the 7.5% subsidy policy, endogenous entry is less important: it has a negligible impact on surplus, and increases revenue by only about 1% increase relative to an unsubsidized auction.

Are there any weaknesses of subsidy policies? One is that the subsidy level must be carefully chosen, but our results suggest that a wide range of subsidy levels improve over current policy. Another concern is that setting aside a certain fraction of timber *guarantees* that a minimum amount will be won by small businesses. A subsidy policy does not provide a firm guarantee. This concern may be particularly salient in a one-off auction setting, such as a sale of radio spectrum. But with many similar sales, it may be less important. To assess this, we used the model to compute the probability distribution of the timber volume won by small businesses in the 1330 sales in our data under the observed set-aside policy and under a 7.5% subsidy. The cumulative distributions functions are shown in Figure 1. Because of the guarantee that set-asides provide, the 7.5% subsidy CDF cannot quite stochastically

dominate the set-aside CDF. Nevertheless, Figure 1 shows that the relation between the CDFs is very nearly one of stochastic dominance. It is very unlikely that under this subsidy loggers win less than the quantity guaranteed to them by set-asides. Considering features of outcome distributions other than the means does not much alter the basic picture: set-asides appear to be a relatively expensive way to achieve distributional goals.

6. Subsidies and Optimal Auction Design

Our results on subsidies can be connected usefully to the theory of optimal auction design. The connection is easiest to see with fixed participation. Suppose that a sale attracts some combination of big and small bidders, N in total. Recall that if bidder i is of type $\tau \in \{S, B\}$ and has value v , its “marginal revenue” is $MR_i(v) = v - \frac{1-F_\tau(v)}{f_\tau(v)}$. Let $\mathbf{v} = (v_1, \dots, v_N)$ denote the vector of bidder values, and let $q_i(\mathbf{v})$ denote the equilibrium probability that i wins as a function of the values. (Generally q_i will be zero or one, unless there are ties or a random allocation.) Of course, q_i may depend on the size of the bidders, the auction format, and whether a subsidy is in place.

Standard results from auction theory relate the allocation rule q_1, \dots, q_N and the marginal revenue functions MR_1, \dots, MR_N to the expected surplus and revenue from the auction. In particular, expected auction surplus is $\mathbb{E}_{\mathbf{v}} [\sum_i q_i(\mathbf{v}) \cdot v_i]$, while expected revenue is $\mathbb{E}_{\mathbf{v}} [\sum_i q_i(\mathbf{v}) \cdot MR_i(v_i)]$. In general, shifting the allocation toward bidders with higher value increases expected surplus, while shifting the allocation toward bidders with higher marginal revenue increases expected revenue. An efficient auction awards the sale to the bidder with the highest value — so $q_i = 1$ if and only if $v_i \geq v_j$ for all j . A revenue optimal auction awards the sale to bidder i if and only if $MR_i(v_i) \geq MR_j(v_j)$ for all j , and $MR_i(v_i) \geq 0$.

Figure 2 represents different allocations in the space of bidder valuations, assuming a tract with average characteristics ($X = \bar{X}$, $u = 1$). The x-axis represents the highest small bidder valuation, and the y-axis the highest big bidder valuation. The forty-five degree line represents the efficient allocation: the high-value small bidder should win if and only if its value is greater than that of the high-value big bidder. The left-most curve describes the revenue-maximizing allocation, which favors small bidders. For all points to the right of the

curve, the high-value small firm has the highest marginal revenue, so shifting the allocation from the forty-five degree line toward the revenue-maximizing allocation reduces efficiency but increases revenue.

The remaining curves in Figure 2 describe the equilibrium allocations from open and sealed bid auctions with no subsidy and with a 7.5% small bidder subsidy. (The open auction allocations do not depend on the number of bidders; the sealed allocations assume two small and two big bidders.) An open auction with no subsidy yields an efficient outcome. Both a shift to sealed bidding and a small bidder subsidy shift the allocation toward small bidders. Both changes increase revenue at some cost to efficiency. To first order, however, a small shift away from the efficient allocation matters more for revenue, helping to explain why the revenue boost from a subsidy dominates the efficiency loss in Table 7.

Figure 2 also shows that in the range of likely valuations (the deciles of the high bidder value distributions are plotted on the axes), the subsidy has a larger effect on the allocation than moving from open auctions to sealed bidding. This helps to explain why ALS found relatively minor effects of shifting between open and sealed bidding under competitive bidding, while we find larger effects from modest subsidies. Of course, a set-aside policy has even more dramatic consequences because it shifts the allocation to coincide with the y-axis — reducing both efficiency and revenue if small firm participation is held constant.

A final point in Figure 2 concerns endogenous participation. Suppose that as in the Figure, we start from a situation where equilibrium involves both types of bidders entering. Because small bidders make lower expected profits conditional on entry, shifting the allocation in their favor will increase small bidder participation without decreasing the big bidder participation, at least for small changes in the allocation. Now, in the unsubsidized open auction case, equilibrium entry is efficient. So if we shift to the small bidder subsidy case, it follows that endogenous entry will tend to reinforce both the increase in revenue and the reduction in surplus. This helps to clarify the findings reported in Table 8.

7. Conclusion

Distributional objectives are an important feature of public sector procurement and natural resource sales. They can be achieved in a variety of ways, with subsidies and set-asides

being perhaps the two most common. Economic theory is not dispositive on which approach can achieve a given distributional goal at lower social cost. Our estimates from the federal government's timber sales program, however, provide an example where set-asides might in practice be relatively costly compared to a subsidy policy. The logic underlying our results is that if the goal is to favor a significantly weaker set of bidders, it may be better to subsidize the weaker bidders and modestly tip outcomes across a broad range of sales, rather than setting aside a targeted number of sales and precluding efficient firms from entering. Of course, a qualification of our results is that they are obtained from a relatively small data sample and a particular federal program. It would be interesting to explore whether there are larger classes of public sector procurement or resource sale problems where similar results obtain.

References

- Andrews, Donald W.K.**, “A Conditional Kolmogorov Test,” *Econometrica*, 1997, 65(5), pp. 1097-1128.
- Athey, S.**, “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” *Econometrica*, 2001, 69(4), pp. 861-889.
- Athey, S. and P. Haile**, “Identification in Standard Auction Models,” *Econometrica*, 2002, 70 (6), pp. 2107-2140.
- Athey, S. and J. Levin**, “Information and Competition in U.S. Forest Service Timber Auctions,” *Journal of Political Economy*, 2001, 109(2), pp. 375–417.
- Athey, S., J. Levin and E. Seira**, “Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions,” *Quarterly Journal of Economics*, 2011, 126(1).
- Ayres, I. and P. Cramton**, “Deficit Reduction through Diversity: How Affirmative Action at the FCC Increased Competition,” *Stanford Law Review*, 1996, 48(4), pp. 761–815.
- Bajari, P.**, “Comparing Competition and Collusion: A Numerical Approach,” *Economic Theory*, 2001, 18(1), pp. 187-205.
- Baldwin, L., R. Marshall, and J.-F. Richard.**, “Bidder Collusion at U.S. Forest Service Timber Auctions,” *Journal of Political Economy*, August 1997, 105(4), pp. 657–699.
- Branman, L., and L. Froeb**, “Mergers, Cartels, Set-Asides and Bidding Preferences in Asymmetric Second-price Auctions,” *Review of Economics and Statistics*, 2000, 82(2), pp. 283-290.
- Bulow, J. and P. Klemperer**, “Prices and the Winner’s Curse,” *Rand Journal of Economics*, 2001, 33(1), pp. 1-21.
- Flambard, V. and I. Perrigne**, “Asymmetry in Procurement Auctions: Evidence from Snow Removal Contracts,” *Economic Journal*, 2006, 116(514), pp. 1014-1036.
- Guerre, E., I. Perrigne and Q. Vuong**, “Optimal Nonparametric Estimation of First Price Auctions,” *Econometrica*, May 2000, 68(3), pp. 525–574.
- Froeb, L. and P. McAfee**, “Deterring Bid Rigging in Forest Service Timber Auctions,” U.S. Department of Justice Economic Analysis Group Working Paper, 1988.
- Haile, P.**, “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales,” *American Economic Review*, June 2001, 92(3), pp. 399–427.

- Haile, P., and E. Tamer**, “Inference with an Incomplete Model of English Auctions,” *Journal of Political Economy*, 2003, 111(1), pp. 1-51.
- Krasnokutskaya, E.**, “Identification and Estimation of Auction Models with Unobserved Heterogeneity,” *Review of Economic Studies*, forthcoming.
- Krasnokutskaya, E. and K. Seim**, “Bid Preference Programs and Participation in Highway Procurement Auctions,” *American Economic Review*, forthcoming.
- Li, T.** “Econometrics of First-Price Auctions with Entry and Binding Reserve Prices,” *Journal of Econometrics*, 2005, 126(1), pp. 173-200.
- Li, T. and X. Zheng**, “Entry and Competition in First-Price Auctions: Theory and Evidence from Procurement Auctions,” *Review of Economic Studies*, 76, pp. 1397-1429.
- Marion, J.** “Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions,” *Journal of Public Economics*, 2007, 91(7-8), pp. 1591-1624.
- Marmer, V., A. Shneyerov and P. Xu**, “What Model for Entry in First-Price Auctions: A Nonparametric Approach,” University of British Columbia Working Paper, 2010.
- Marshall, R., M. Meurer, J.F. Richard and W. Stromquist**, “Numerical Analysis of Asymmetric First-Price Auctions,” *Games and Economic Behavior*, 1994, 7, pp. 193-220.
- Maskin, E. and J. Riley**, “Asymmetric Auctions,” *Review of Economic Studies*, July 2000, 67(3), pp. 413-438.
- McAfee, P. and J. McMillan**, “Government Procurement and International Trade,” *Journal of International Economics*, 1989, 26, pp. 291-308.
- Myerson, R.**, “Optimal Auction Design,” *Mathematics of Operations Research*, 1981, 6, pp. 58-73.
- Roberts, J. and A. Sweeting**, “Entry and Selection in Auctions,” Duke University Working Paper, 2010.

Table 1: Summary Statistics

N	All Sales				Open Sales				Sealed Sales			
	Unrestricted		Set-Aside		Unrestricted		Set-Aside		Unrestricted		Set-Aside	
	1167		163		786		127		381		36	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Auction Outcomes												
Prices (\$/mbf)	95.12	69.05	90.74	58.43	97.12	70.20	94.55	62.32	90.99	66.51	77.30	39.78
Entrants	4.26	2.50	4.50	2.31	4.33	2.39	4.52	2.30	4.13	2.69	4.42	2.39
# Small Bidders Entering	2.71	2.19	4.50	2.31	2.29	1.89	4.52	2.30	3.59	2.48	4.42	2.39
# Big Bidders Entering	1.55	1.44	0.00	0.00	2.04	1.38	0.00	0.00	0.54	0.95	0.00	0.00
Small Bidder Wins Auction	0.52	0.50	1.00	0.00	0.37	0.48	1.00	0.00	0.85	0.36	1.00	0.00
Appraisal Variables												
Volume of timber (hundred mbf)	51.23	50.95	52.96	45.25	71.18	49.89	64.75	43.73	10.07	17.32	11.37	17.59
Small	0.21	0.41	0.14	0.35	0.04	0.21	0.03	0.18	0.54	0.50	0.53	0.51
Reserve Price (\$/mbf)	40.33	35.02	37.15	28.20	40.01	34.00	37.85	30.20	41.00	37.07	34.70	19.68
Selling Value (\$/mbf)	272.88	97.56	292.05	66.27	281.79	85.17	301.60	55.35	254.48	117.10	258.36	88.33
Road Construction (\$/mbf)	7.24	12.04	7.96	11.17	10.16	13.37	10.02	11.83	1.21	4.62	0.71	2.07
No Road Construction	0.01	0.08	0.40	0.49	0.01	0.08	0.02	0.12	0.01	0.07	0.00	0.00
Logging Costs (\$/mbf)	102.87	40.99	109.94	32.93	107.22	36.50	114.06	30.59	93.88	47.79	95.40	37.03
Manufacturing Costs (\$/mbf)	121.97	46.01	127.76	34.76	127.55	38.08	132.42	26.39	110.45	57.47	111.33	52.25
Sale Characteristics												
Contract Length (months)	24.25	14.95	27.87	14.26	30.32	13.75	31.51	13.55	11.72	7.86	15.00	7.87
Species Herfindal	0.57	0.24	0.59	0.26	0.57	0.24	0.59	0.26	0.57	0.24	0.58	0.27
Density of Timber (hmbf/acres)	11.60	15.23	11.30	15.04	10.70	14.06	10.06	15.35	13.46	17.27	15.69	13.16
Sealed Bid Sale	0.33	0.47	0.22	0.42	0.00	0.00	0.00	0.00	1.00	0.00	1.00	0.00
Scale Sale	0.81	0.39	0.85	0.36	0.90	0.30	0.92	0.27	0.61	0.49	0.61	0.49
Quarter of Sale	2.42	1.00	2.28	1.04	2.33	1.03	2.21	1.05	2.60	0.90	2.50	0.97
Year of Sale	85.14	2.10	84.98	2.10	85.08	2.08	84.91	2.05	85.25	2.14	85.25	2.26
Housing Starts	1597.42	257.20	1600.46	287.93	1601.92	260.60	1607.35	292.14	1588.14	250.13	1576.17	275.13
Logging companies in county	20.00	18.59	22.11	24.97	21.33	19.39	22.33	24.51	17.24	16.52	21.33	26.89
Sawmills in County	6.33	6.99	8.23	10.99	6.62	7.12	8.25	10.78	5.74	6.69	8.17	11.86
Potential Competition												
Active Small Bidders	13.18	7.44	14.19	7.31	13.31	7.68	14.13	7.20	12.92	6.91	14.39	7.77
Active Big Bidders	3.17	1.76	2.34	1.72	3.34	1.82	2.34	1.79	2.82	1.58	2.36	1.44

Note: Small bidders are those eligible for small business set-aside sales. A "small" tract is one in the lowest quintile by volume. Timber volume is measured in "mbf". Loggers do not have manufacturing capacity, sawmills do. A bidder is "active" if it has participated in any auction in the same forest-district and the same

Table 2: Small and Big Firm Bidding Differences

	Regression Model: Log(Bid)		Sealed Bids Per Quintile	Fraction of Unrestricted Sales Won by Small Firms
	Coefficient	Std Err.		
Big Firm Dummy x				
Sale in First Size Quintile	-0.005	(0.071)	824	0.95
Sale in Second Size Quintile	0.111	(0.040)	752	0.77
Sale in Third Size Quintile	0.041	(0.102)	80	0.51
Sale in Fourth Size Quintile	0.028	(0.138)	46	0.41
Sale in Fifth Size Quintile	0.066	(0.144)	29	0.28
Total Number of Sealed Bids		1731		
Total Number of Auctions		417		1167

Note: The first two columns report regression results where the dependent variable is the logarithm of the bid per unit volume, the data includes all sealed bids submitted in unrestricted auctions, and the explanatory variables include auction fixed effects and a dummy equal to one if the bidder is a big firm interacted with the size of the sale. The sales are assigned to size quintiles based on the volume of timber being sold. The third column shows the number of unrestricted sealed bids for each size category. The fourth column shows the fraction of (sealed and open) unrestricted sales won by small firms.

Table 3: Choice of Set-Aside Sale

	Marginal Effect	Std. Err.
<i>Appraisal Controls</i>		
Ln(Reserve Price)	-0.009	(0.012)
Ln(Selling Value)	0.039	(0.031)
Ln(Manufacturing Costs)	0.012	(0.015)
Ln(Logging Costs)	-0.133	(0.041)
Ln(Road Costs)	0.006	(0.006)
Road Costs Missing (Dummy)	0.111	(0.144)
Appraisal Missing (Dummy)	-0.123	(0.028)
<i>Other Sale Characteristics</i>		
ln(Contract Length/volume)	1.382	(0.576)
Species Herfindal	-0.019	(0.031)
Density of Timber (10,000 mbf/acres)	0.002	(0.045)
Sealed Bid (Dummy)	-0.012	(0.021)
Scale Sale (Dummy)	0.013	(0.020)
ln(Monthly US House Starts)	0.049	(0.080)
<i>Volume Controls (Dummy Variables):</i>		
Volume: 1.5-3 hundred mbf	0.017	(0.061)
Volume: 3-5	0.089	(0.115)
Volume: 5-8	0.062	(0.101)
Volume: 8-12	0.051	(0.102)
Volume: 12-20	0.028	(0.087)
Volume: 20-40	0.168	(0.176)
Volume: 40-65	0.161	(0.165)
Volume: 65-90	0.175	(0.179)
Volume: 90+	0.051	(0.099)
<i>Local Industry Activity</i>		
ln(Loggers in County)	-0.010	(0.013)
ln(Sawmills in County)	-0.006	(0.011)
ln(Active Small Firms)	0.031	(0.012)
ln(Active Big Firms)	-0.047	(0.009)
<i>Additional Controls (Dummy Variables)</i>		
<i>Chi-Squared Statistics (p-value in parenthesis)</i>		
Years	4.760	(0.690)
Quarters	8.620	(0.035)
Species	10.230	(0.115)
Location	37.380	(0.000)
		N=1315
LR chi2 (52)		235.04
P-value		0.00
Pseudo-R2		0.24

Note: Table reports results from a logit regression where the dependent variable is equal to one if the sale is a small business set-aside. The estimates are reported as marginal probability effects at the mean of the independent variables.

Table 4: Effects of Set-Aside Provision

<i>Dependent Variable</i>	(1)		(2)	
	coefficient	s.e.	coefficient	s.e.
<i>Average Effect of Set-Aside</i>				
OLS (full interactions)	-0.127	(0.068)	-0.102	(0.072)

Note: Table reports estimates from OLS regressions of dependent variable on sale characteristics and a dummy variable for small business set-aside sale interacted with the characteristics. The reported estimates are average effects across all tracts. Sale characteristics in the regression include all the variables from Table 3 and the estimated probability that the auction was conducted as a set-aside (the propensity score from the logit regression in Table 3).

Table 5: Bid Distributions for Sealed Bid Auctions

	Panel A: Large Tracts		Panel B: Small Tracts	
	Coefficient	Std. Err.	Coefficient	Std. Err.
	$\ln(\lambda)$		$\ln(\lambda)$	
Ln(Reserve Price)	0.600	(0.050)	0.597	0.047
Ln(Selling Value)	-0.119	(0.045)	-0.096	0.041
Ln(Manufacturing Costs)	0.151	(0.083)	0.029	0.031
Ln(Logging Costs)	0.164	(0.177)	-0.310	0.176
Appraisal Missing (Dummy)	0.537	(0.813)	-1.647	0.799
Ln(Road Costs)	-0.034	(0.029)	0.156	0.069
Road Costs Missing (Dummy)	-0.005	(0.227)	-	-
Species Herfindal	-0.227	(0.105)	-0.265	0.121
Density of Timber (10,000 mbf/acres)	-0.154	(0.160)	0.007	0.192
Scale Sale (Dummy)	0.155	(0.064)	-0.027	0.058
Big Bidder (Dummy)	0.107	(0.039)	-	-
Big Bidder * Big Bidder Entrants=1	0.017	(0.064)	-	-
Min(Small Bidder Entrants,5)	0.094	(0.022)	-	-
Min(Big Bidder Entrants,5)	0.063	(0.025)	-	-
Min(Total Entrants,10)	-	-	0.069	0.011
Volume - 1st Decile	-	-	0.031	0.054
Volume - 3rd Decile	0.054	0.300	-	-
Volume - 4th Decile	0.005	0.298	-	-
Volume - 5th Decile	-0.033	0.296	-	-
Volume - 6th Decile	-0.403	0.317	-	-
Volume - 7th Decile	-0.087	0.317	-	-
Volume - 8th Decile	-0.152	0.316	-	-
Volume - 9th Decile	-0.070	0.324	-	-
Constant	0.661	1.003	3.851	0.978
Additional Controls	Forest and Year Dummies		Forest and Year Dummies	
	$\ln(\rho)$		$\ln(\rho)$	
Big Bidder	-0.163	(0.093)	-	-
Big Bidder * Big Bidder Entrants=1	0.287	(0.165)	-	-
Min(Small Bidder Entrants,5)	0.064	(0.025)	-	-
Min(Big Bidder Entrants,5)	0.032	(0.027)	-	-
Min(Total Entrants,10)	-	-	0.008	0.013
Constant	0.987	(0.118)	1.175	0.084
	$\ln(\theta)$		$\ln(\theta)$	
Constant	-1.121	(0.233)	-0.603	(0.171)
	N = 797		N = 710	
	Wald χ^2 (42)	1596.22	Wald χ^2 (31)	624.53
	Prob > chi2	0.000	Prob > chi2	0.00

Note: Table presents the maximum likelihood estimates of the Gamma-Weibull bidding model, described in the text, run separately on large and small unrestricted entry sealed bid auctions with 2 or more bidders. An observation is a sealed bid. Scale controls also include dummy variables equal to one if sale had no road construction, or if appraisal variables are missing.

Table 6: Effect of the Set-Aside Program

	Actual	Unrestricted	Predicted Outcomes		
	Outcome		Std. Err.	Set-Aside	Std. Err.
<i>Panel A: Tracts sold by unrestricted sale (N=1167)</i>					
Avg. Small Bidder Entry	2.71	2.71	(0.00)	4.64	(0.62)
Avg. Big Bidder Entry	1.55	1.55	(0.00)	0.00	(0.00)
Avg. Total Entry	4.26	4.26	(0.00)	4.64	(0.62)
Avg. Prices	95.12	98.27	(7.79)	85.40	(7.61)
% Sales won by Small Bidders	52.44	54.36	(0.02)	97.48	(0.01)
Avg. Sale Surplus (per mbf)		136.68	(34.38)	107.80	(23.02)
<i>Panel B: Tracts sold by set-aside sale (N=163)</i>					
Avg. Small Bidder Entry	4.50	3.50	(0.13)	4.66	(0.57)
Avg. Big Bidder Entry	0.00	1.66	(0.00)	0.00	(0.00)
Avg. Total Entry	4.50	5.16	(0.13)	4.66	(0.57)
Avg. Prices	90.74	106.24	(8.62)	89.18	(8.46)
% Sales won by Small Bidders	100.00	60.25	(0.03)	99.04	(0.00)
Avg. Sale Surplus (per mbf)		139.08	(39.72)	108.65	(29.34)

Note: Table reports actual outcomes for California sales and predicted outcomes generated by the model, assuming either that sales are unrestricted or conducted as set-asides. The predicted fraction of set-asides won by small bidders can be less than one because (a) for some unrestricted sales, there are no potential small entrants, and (b) the reserve price binds in some simulated auctions. The possibility for multiple entry equilibria means we obtain a range of predicted outcomes. We report the midpoints; the ranges are less than 0.5% around the reported numbers. Bootstrapped standard errors from 100 bootstrap repetitions are reported in parentheses.

Table 7: Comparison of Set-Asides and Subsidies

<i>Outcome Variable</i>	Bidder Entry		Bidder Profits (\$000s)		% Sales Won	% Vol. Won	Price (\$/mbf)	Surplus (\$/mbf)
	Small Bidders	Big Bidders	Small Bidders	Big Bidders	by Small Bidders	by Small Bidders		
<i>Panel A: All Sales (N=1330)</i>								
Actual Policy	2.95	1.36	76.1	221.0	59.8	42.9	97.2	133.2
Volume Set-Asides	3.26	1.05	77.7	218.9	67.9	42.9	93.6	126.7
"Optimal" Set-Asides	2.96	1.48	68.5	235.7	58.0	43.0	99.3	136.5
No Subsidy	2.81	1.57	63.9	245.0	55.1	36.8	99.3	137.0
2.5% Subsidy	3.15	1.56	70.5	232.9	58.6	42.0	100.9	136.9
5% Subsidy	3.19	1.56	73.5	227.8	59.8	43.6	101.3	136.9
7.5% Subsidy	3.26	1.56	76.7	222.3	61.0	45.4	101.7	136.8
10% Subsidy	3.34	1.56	79.9	216.7	62.4	47.2	102.2	136.7
<i>Panel B: Small Sales (N=265)</i>								
Actual Policy	3.63	0.18	11.61	1.07	93.7	92.6	82.91	107.55
Volume Set-Asides	3.64	0.01	12.24	0.25	97.0	97.2	81.31	106.30
"Optimal" Set-Asides	3.63	0.19	11.58	1.11	93.3	92.2	83.00	107.61
No Subsidy	3.63	0.19	11.58	1.11	93.3	92.2	83.00	107.61
2.5% Subsidy	3.65	0.17	11.65	1.04	94.0	93.0	83.01	107.60
5% Subsidy	3.65	0.16	11.68	0.99	94.3	93.4	83.01	107.60
7.5% Subsidy	3.65	0.16	11.71	0.96	94.4	93.6	83.00	107.59
10% Subsidy	3.65	0.15	11.74	0.93	94.6	93.8	82.95	107.56
<i>Panel C: Large Sales (N=1065)</i>								
Actual Policy	2.78	1.66	92.1	275.7	51.4	42.3	100.7	139.6
Volume Set-Asides	3.16	1.31	94.0	273.3	60.6	42.3	96.6	131.8
"Optimal" Set-Asides	2.80	1.80	82.7	294.1	49.2	42.4	103.4	143.7
No Subsidy	2.60	1.91	77.0	305.7	45.6	36.1	103.3	144.3
2.5% Subsidy	3.02	1.91	85.2	290.6	49.8	41.5	105.4	144.2
5% Subsidy	3.08	1.91	88.9	284.2	51.2	43.0	105.9	144.2
7.5% Subsidy	3.17	1.91	92.9	277.4	52.7	44.8	106.4	144.1
10% Subsidy	3.27	1.91	96.9	270.4	54.3	46.7	106.9	143.9

Note: Table reports outcomes, averaged over all auctions, of different set-aside and subsidy policies. The actual outcomes are those given the set-aside policy observed in the data. All set-aside policies shown achieve at least as much total small bidder volume over all sales as actual set-asides do. "Volume Set-Asides" set aside those sales smaller than a volume threshold, where the volume threshold is set to achieve the overall small bidder volume target. "Optimal Set-Asides" are those sales which have the fitted values closest to zero in a regression of (Unrestricted Surplus - Set-Aside Surplus)/(Unrestricted Small Bidder Volume - Set-Aside Small Bidder Volume) on tract characteristics. 77 sales are set aside under optimal set-asides, 415 under volume-based set-asides, and 66 under predicted optimal set-asides. For the counterfactual subsidies, the subsidies are applied to every auction in the sample.

Table 8: Decomposing Bidding and Participation Effects

<i>Outcome Variable</i>	Small Bidder Entry	Big Bidder Entry	Total Small Bidder Prof. (\$'000s)	Total Big Bidder Prof. (\$'000s)	% Sales Won by Small Bidder	% Volume Won by Small Bidder	Price (\$/mbf)	Surplus (\$/mbf)
<i>Panel A: Unrestricted Sales with Positive Small Bidder Entry (N=1058)</i>								
No Subsidy or Set-Aside	2.99	1.54	67.5	230.3	60.0	41.3	102.6	140.3
Set-Aside, Fixed Small Entry	2.99		202.9		96.4	95.4	76.0	104.9
Set-Aside, Endogenous Entry	4.64		190.7		97.4	97.0	86.3	109.2
7.5 % Subsidy, Fixed Big, Small Entry	2.99	1.54	74.9	217.7	62.9	45.3	103.3	140.1
7.5 % Subsidy, Endogenous Entry	3.50	1.53	80.8	206.8	66.1	50.5	105.3	140.1
<i>Panel B: Set-Aside Sales (N=163)</i>								
No Subsidy or Set-Aside	3.50	1.66	83.7	195.9	60.2	50.2	106.2	139.1
Set-Aside, Fixed Small Entry	3.50		176.0		91.3	87.7	77.8	100.1
Set-Aside, Endogenous Entry	4.66		182.6		99.0	98.9	89.2	108.6
7.5 % Subsidy, Fixed Big, Small Entry	3.50	1.66	89.8	185.0	63.0	53.4	106.9	138.9
7.5 % Subsidy, Endogenous Entry	3.74	1.65	93.2	179.1	64.9	55.7	108.0	138.9

Note: Table reports predicted outcomes from a set-aside and from a 7.5% subsidy, assuming that entry is either fixed or endogenous. For the set-aside rows, "Fixed Small Entry" refers to fixing small bidder entry at its equilibrium value in the unrestricted, unsubsidized case, with big bidder entry set to zero. For the 7.5% subsidy rows, "Fixed Big, Small Entry" refers to fixing big and small bidder entry at their equilibrium values in the unrestricted, unsubsidized case. "Endogenous entry" means equilibrium entry under a set-aside or a 7.5% subsidy. Also included are "No Subsidy or Set-Aside" outcomes. These are the model predictions assuming unrestricted equilibrium entry and unsubsidized sales. Reported outcomes are averages across sales. Panel A drops the 109 unrestricted sales for which there are no small entrants, since these sales would fail in the "set-aside, fixed entry" counterfactual. Note that small bidders do not win all set-aside sales in the simulations because the reserve price occasionally binds.

Figure 1: Empirical CDFs




Figure 2: Winning Regions for Big and Small Bidders, Open and Sealed Auctions

