

NBER WORKING PAPER SERIES

SET-ASIDES AND SUBSIDIES IN AUCTIONS

Susan Athey
Dominic Coey
Jonathan Levin

Working Paper 16851
<http://www.nber.org/papers/w16851>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2011

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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Susan Athey, Dominic Coey, and Jonathan Levin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Set-Asides and Subsidies in Auctions
Susan Athey, Dominic Coey, and Jonathan Levin
NBER Working Paper No. 16851
March 2011
JEL No. D44,H57,L53

ABSTRACT

Set-asides and subsidies are used extensively in government procurement and natural resource sales. We analyze these policies in an empirical model of U.S. Forest Service timber auctions. The model fits the data well both within the sample of unrestricted sales where we estimate the model, and when we predict (out of sample) bidder entry and prices for small business set-asides. Our estimates suggest that restricting entry to small businesses substantially reduces efficiency and revenue, although it does increase small business participation. An alternative policy of subsidizing small bidders would increase revenue and small bidder profit, while eliminating almost all of the efficiency loss of set-asides, and only slightly decreasing the profit of larger firms. We explain these findings by connecting to the theory of optimal auction design.

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

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

Dominic Coey
Stanford University
Department of Economics
579 Serra Mall
Stanford CA 94305-6072
dcoey@stanford.edu

An online appendix is available at:
<http://www.nber.org/data-appendix/w16851>

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 an econometric 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 setting 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 17% and cost the Forest Service about 5% in revenue. Providing a subsidy to small bidders in all auctions appears to be a more effective means of achieving distributional goals. A range of subsidies might have eliminated both efficiency and revenue losses, while allocating the same volume of timber to small bidders, increasing aggregate small firm profits, and

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.

only slightly reducing the profit of larger firms. If other U.S. procurement and resource allocation programs 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 more large firms to begin with, 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

³Suppose, for instance, there are two large firms. Then absent an entry restriction, both large firms enter in equilibrium, giving an expected price of 10. And a participation restriction decreases both revenue and social surplus. The effects of restricting participation also can depend on features other than costly participation, e.g. Bulow and Klemperer (2001) show that restricting participation sometimes can be beneficial if there are strong "winner's curse" effects.

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.

We next develop 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 (i.e. varying entry restrictions and subsidy levels), it is important to have an econometric model that can accurately predict entry and prices “out-of-sample” as 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 5% of observed values, and we cannot reject equality of the predicted and observed bid distributions.

One observation from comparing unrestricted and set-aside sales is 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 (30 and 28 percent, respectively, rather than 5 and 17 percent).

We also use the model to calculate the 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 of subsidies

⁴The idea that the Forest Service set-aside program could be replaced with a subsidy policy is discussed

seem more effective at achieving distributional goals than the observed policy of set-asides. From a programmatic perspective, a 6% subsidy for small businesses would result in small firms winning as much timber as under the set-aside program, with 4% higher prices and a 2% increase in overall program efficiency. There is a small decline in the expected profit of larger firm (less than 2%), which disappears entirely with a slightly smaller subsidy. The attractive performance of subsidies relative to set-asides 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 limit efficiency losses while achieving distributional objectives is to select sales to be set-asides where efficiency losses would be small. We construct a statistical 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%.

by Froeb and McAfee (1988), and also by Brannman and Froeb (2000).

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.⁶ 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. With sealed bidding, there is a unique Bayesian Nash equilibrium in which the bidders bid their values minus a shading factor that depends on the equilibrium behavior of opponents. We state and use the first order conditions for equilibrium in Section 4. With an open auction, there is an equilibrium in

⁶There are also a number of specific modeling differences, for instance in the way the papers model entry behavior (mixed strategies versus pure strategies with incomplete information about entry costs), what bidders know about their competitors when they submit their bids, and in the parametric modeling of bid distributions.

weakly dominant strategies in which each bidder continues in the auction until the price reaches its valuation, at which point it drops out.

In the entry game, we focus on type-symmetric equilibria. Small and big bidders enter with probabilities denoted (p_S, p_B) , and entrants earn an expected profit of at least K . In set-asides, $p_B = 0$. Later, we distinguish three types of non-set-aside sales, based on our observation that large firms appear to have higher values than small firms in most sales, but not the very smallest sales. In typical sales, where we estimate that large firms have substantially higher values, we focus on equilibria in which big bidders enter the auction ($p_B = 1$), and small bidders randomize their entry with equal probability.⁷ In small sales without subsidies, $F_S = F_B$ and we focus on the unique fully symmetric equilibrium in which $p_S = p_B$. In small sales with a subsidy for small bidders, we consider the full set of type-symmetric equilibria.

We should emphasize that modeling entry requires a number of choices that can be debated. For instance, our focus on type-symmetric entry equilibria involves looking at mixed strategies. While symmetry is a standard restriction, mixed strategy equilibria have the somewhat unintuitive property that decreasing the number of potential bidders — for example, due to a set-aside program — potentially can increase expected participation and revenue.⁸ This “coordination effect” originally motivated us to focus on pure strategy entry equilibria. But as we discuss in Section 4, that approach led us to estimate implausibly high entry costs for many low entry sales. Subsequently we found that our estimates with mixed strategy entry do not imply significant coordination effects, alleviating our initial concern.

In addition to deciding 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

⁷With a sufficiently large number of big bidders, another possibility would be that big bidders randomize entry and small bidders for sure do not enter. As discussed in more detail in ALS, this does not appear to be the empirically relevant case in our setting.

⁸This can happen even if all bidders are symmetric. Consider 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 $10/3$.

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). 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.⁹ 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

⁹This 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.

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 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.¹⁰ 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

¹⁰See the U.S. Forest Service Handbook, Section 2409.18 on Timber Sale Preparation.

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 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 distinguish 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.¹¹ 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 complication to the already challenging problem of accurate computation. We provide some 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

¹¹We 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.

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 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 A3),

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 both over all tracts in the sample, and over those tracts that were actually sold as set-asides (this is $\hat{\delta} + \bar{X}\hat{\gamma}$, where \bar{X} is the vector of the mean covariates of the sales in question). The estimates suggest that set-asides reduce revenue and entry on an average tract, and may increase entry on an average set-aside tract. 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. From these value distributions, we obtain bidder profits conditional on entry. We then estimate equilibrium entry probabilities for each auction if entry of large bidders were unrestricted, as a function of sale characteristics and entry patterns in the local geographic area. These entry probabilities, together with profits and our measure of potential entry, allow us to infer entry costs. 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 Guerre, Perrigne and Vuong (2000), Krasnokutskaya (2011) and ALS. 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.¹² 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

¹²Krasnokutskaya 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.

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 n . We assume that conditional on (X, u, n) , the bids of small and big firms have a Weibull distribution.¹³ 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 5 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

¹³Athey, 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.

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:

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

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 14.2% 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 to 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 bidder is 2.94 times the expected profit of a small bidder; the ratio is 2.83 in a sealed auction.

C. Potential Entrants, Entry Probabilities and Entry Cost

We next estimate the bidder entry costs for each sale, using the assumption that observed entry patterns follow a type-symmetric equilibrium. To do this, we first need to estimate the number of potential small and big entrants at each sale. We assume that potential small bidder entry at a given sale is equal to the maximum small bidder entry observed across all the unrestricted sales in the same forest-year. We can do somewhat better for the big bidders because for the larger sales in our sample, the asymmetry of the estimated value distributions implies that if small bidders are entering with positive probability, as we observe, then big bidders must be entering with probability one. So for larger unrestricted sales, we set the number of potential big bidders equal to the number of actual big entrants. Then for the remaining sales, we assume that the number of potential big bidders equals the number of actual big entrants in the “most similar” larger unrestricted sale, where “most similar” means the sale in the same forest-year that is closest in timber volume.

To estimate entry costs, we use the equilibrium entry condition for small bidders. At a type-symmetric equilibrium with $0 < p_S < 1$, these bidders are just indifferent between entering and not entering. Denoting potential entry as $N = (N_S, N_B)$, and writing the entry cost as $K(X)$, the indifference condition is:

$$\sum_{n \subset N} \pi_S(X, n) Pr[n|X, N, i \in n] = K(X). \quad (4)$$

Here $Pr[n|X, N, i \in n]$ is the probability that $n = (n_S, n_B)$ bidders enter given that i enters. Assuming i is a small bidder,

$$Pr[n|X, N, i \in n] = Pr[n_B|X, N] \cdot Pr[n_S|X, N, i \in n] \quad (5)$$

$$= Pr[n_B|X, N] \cdot \binom{N_S - 1}{n_S - 1} p_S^{n_S - 1} (1 - p_S)^{N_S - n_S}. \quad (6)$$

That is, the number of small entrants follows a binomial distribution where each of i 's opponent's independently decides to enter with probability p_S . The big entrants either enter with probability one (so $n_B = N_B$) if the sale is large and the value distributions are asymmetric, or enter independently with probability $p_B = p_S$ if the sale is small and the value distributions are symmetric.

We estimate the equilibrium entry probability p_S using the observed data on small bidder entry, using the following parametric model:

$$p_S(X, N) = \frac{\exp(X\alpha_X + N\alpha_N)}{1 + \exp(X\alpha_X + N\alpha_N)}. \quad (7)$$

Table 6 reports the maximum likelihood estimates of the entry probability parameters. Using our estimates of $p_S(X, N)$, combined with the estimated profit function $\pi_S(X, n)$, we use equation (4) to infer entry costs for all sales as a function of their covariates (X, N) . We estimate entry costs that are relatively high, \$15,754 for the median sale.

As noted above, our empirical model of entry assumes a type-symmetric equilibrium with mixed strategy entry by small bidders. In an earlier version of the paper, we based our estimation on the assumption of a pure strategy entry equilibrium. The difficulty with that approach is that we identified many sales where it appeared that a single additional small bidder entrant could make substantial profit by entering unless entry costs were very high. In the context of a pure strategy equilibrium, a natural way to rationalize this might be to assume that the set of potential entrants in such a sale had been exhausted. We found it difficult to get the data to separately distinguish high entry costs from a smaller number of potential entrants, and hence hard to avoid estimating entry costs that seemed implausible without making assumptions about the set of potential entrants that substantially affected

our results. The mixed equilibrium model, despite having its own potential shortcomings, avoids some of these problems because it assigns positive probability to sales having ex post the feature that an extra small bidder could enter and make a positive profit.

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.

The results are reported in Table 7, 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 and the percentage of sales won by small bidders are 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

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

bottom panel. The model predicts entry within 5% of the actual value, and prices and the percentage of sales won by small bidders (rather than left unsold) within 2%. We also test whether the predicted sealed bid distribution from the model matches to actual distribution of sealed bids. To do this, we use the model to simulate the set-aside sales, then compare the deciles of the simulated distribution of bids to the deciles of the observed bid distribution. We construct standard errors for this test statistic using a bootstrap method.¹⁵ The p-value associated with this test is 0.52, and the null hypothesis of equality of deciles is not rejected.

B. Assessing the Impact of Set-Asides

Table 7 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 1.71 big bidders per sale, with 1.93 fewer small entrants. The fraction of sales won by small bidders would have dropped to 51%, but revenue would have been 6% higher, and surplus 21% 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 have almost fully compensated in terms of revenue, leading to only a 3% revenue loss. The predicted surplus loss is greater however, at 16%. These results are driven mostly by the improved efficiency in allocation. The reduced entry costs from unrestricted sales contribute around 16% of the gain in surplus.

As previously noted, decreasing the number of potential bidders can in principle increase participation and revenue when bidders play mixed strategies. One way to assess the quantitative importance of this effect is to reduce the number of potential small bidder entrants by one and recompute the entry equilibria and outcomes. If a significant coordination effect is present, this should increase prices or entry. This is not what we find: overall small bidder entry fell as a result of removing a potential small bidder, from 2.65 to 2.59. Prices also fell, from 94.90 to 94.50. The coordination effect does not appear to be very important in our estimated model.

¹⁵To construct a variance-covariance matrix for the observed bid deciles, we resample from the set of set-aside sales and use the bid distributions from each resampled dataset. For the predicted sales, we resample from the asymptotic distribution of the parameter estimates, and for each draw, re-do the simulation of the set-aside sales.

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 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.

We then compute entry equilibria. In the smallest sales, there may be multiple type-symmetric equilibria. Absent subsidies, small and big bidders have identical valuation distributions. We consider the natural symmetric equilibrium in which all potential bidders enter with the same probability. With small bidder subsidies, we compute all type-symmetric entry equilibria, and report average outcomes over these equilibria.

Table 8 shows expected outcomes, averaged over all auctions in the sample, and then with the smallest sales broken out, for 6 different subsidy levels (no subsidy or $\alpha = 5, 6, 10, 15$ and 20%). 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 6%, 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 20% 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 6% subsidy would increase average small business participation in the larger sales from 2.82 to 3.04.

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 6% subsidy entails more total surplus than the observed set-aside policy and almost no surplus loss relative to the no-subsidy, no set-aside case. This subsidy level also results in slightly more sales and timber won by small firms, greater small bidder profits, greater revenue, and almost as great big bidder profits as the set-aside policy. Larger subsidies reduce efficiency further, 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 $(\text{Unrestricted Surplus} - \text{Set-Aside Surplus}) / (\text{Unrestricted Small Bidder Volume} - \text{Set-Aside Small Bidder Volume})$ on tract characteristics, selecting just enough sales so that the volume constraint is satisfied. This approach leads to outcomes almost as efficient as subsidies. Revenue is approximately the same as if there were no preference program at all. Compared to the 6% subsidy, this set-aside policy results in 3% less revenue, 11% less small bidder profit, and 8% more big bidder profit.

Table 9 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 17%, and it reduced the revenue losses from 30% to 5%. For fixed entry, a 5% subsidy has much less dramatic effects on price and surplus than a set-aside. Endogenous entry is thus less important in subsidized sales than set-asides, although it is responsible for most of the change in revenue: The 5% subsidy policy increases revenue over the unsubsidized case, due in part from shifting the allocation to the small bidders conditional on entry, but mainly from encouraging more small bidder entry.

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 6% subsidy. The cumulative distributions functions are shown in Figure 1. Because of the guarantee that set-asides provide, the 6% 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 $MR_i(v) = v - \frac{1-F_r(v)}{f_r(v)}$. 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 6% 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 8.

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 concerns endogenous participation. Suppose that we start from a situation where equilibrium involves the big bidders entering with probability one and the small bidders mixing. 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. Figure 2 also show the revenue-optimal allocation rule with endogenous entry, which involves a larger bias toward small bidders than the 6% subsidy.

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*, 2011, 78(1), pp. 293-327.
- 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*, 2009, 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.
<i>Auction Outcomes</i>												
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 Firms 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 Firms 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 Firm 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
<i>Appraisal Variables</i>												
Volume of timber (100 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 Sale Dummy	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
Road Costs Missing	0.01	0.08	0.40	0.49	0.01	0.08	0.02	0.12	0.01	0.07	0.00	0.00
Appraisal Missing	0.09	0.29	0.02	0.16	0.06	0.23	0.16	0.37	0.01	0.09	0.08	0.28
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
<i>Sale Characteristics</i>												
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 (10,000 mbf/acre)	1160	1523	1130	1504	1070	1406	1006	1535	1346	1727	1569	1316
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	257	1600	288	1602	261	1607	292	1588	250	1576	275
<i>Local Industry Activity</i>												
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
Small Firms Active in Last Year	13.18	7.44	14.19	7.31	13.31	7.68	14.13	7.20	12.92	6.91	14.39	7.77
Big Firms Active in Last Year	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: Data includes non-salvage Forest Service sales in California between 1982 and 1989. Small firms are those that are eligible for small business set-aside sales. A sale is a "small sale" if it is in the lowest quintile by timber volume. Timber volume is measured in thousand board feet (mbf). The measures of local industry activity for a given sale include the number of logging companies and sawmills as reported in the US Census County Business Patterns, and the number of small and big firms that bid in the same forest-district within the prior year.

Table 2: Small and Big Firm Bidding Differences

Regression Model: Log(Bid)				
	Coefficient	Std Err.	Sealed Bids Per Quintile	Fraction of Unrestricted Sealed and Open Sales Won by Small Firms
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)	
	<i>ln(Entrants)</i> coefficient	s.e.	<i>ln(Revenue)</i> coefficient	s.e.
<i>Average Effect (Full Sample)</i>				
OLS with Interactions	-0.127	(0.068)	-0.102	(0.072)
<i>Average Effect on Set-Aside Tracts</i>				
OLS with Interactions	0.083	(0.050)	-0.001	(0.045)

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 all characteristics. The first row shows the estimated average treatment effect of a set-aside for the full sample of sales. The second row shows the estimated effect for the set-aside 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, Species and Year Dummies		Forest, Species 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: Entry Probabilities for Unrestricted Auctions

	Panel A: Large Tracts		Panel B: Small Tracts	
	Coefficient	Std. Err.	Coefficient	Std. Err.
Ln(Reserve Price)	-0.276	(0.054)	-0.354	(0.079)
Ln(Selling Value)	-0.107	(0.079)	0.004	(0.061)
Ln(Manufacturing Costs)	0.143	(0.089)	0.039	(0.064)
Ln(Logging Costs)	-1.637	(0.206)	-1.290	(0.294)
Appraisal Missing (Dummy)	-7.307	(0.966)	-6.252	(1.365)
Ln(Road Costs)	-0.136	(0.027)	-0.102	(0.135)
Road Costs Missing (Dummy)	0.190	(0.379)	-2.304	(1.101)
Species Herfindal	-0.260	(0.146)	-0.339	(0.220)
Density of Timber (10,000 mbf/ac)	-0.332	(0.002)	-0.754	(0.004)
Scale Sale (Dummy)	0.236	(0.103)	0.453	(0.106)
Potential Big Bidder Entrants = 0	0.144	(0.114)	-	-
Min(Potential Big Bidder Entrants)	0.132	(0.035)	-	-
Potential Entrants	-0.190	(0.017)	-0.089	(0.032)
Sealed	0.278	(0.100)	-	-
Volume - 1st Decile	-	-	-0.263	(0.106)
Volume - 3rd Decile	1.242	(0.149)	-	-
Volume - 4th Decile	0.928	(0.137)	-	-
Volume - 5th Decile	0.501	(0.124)	-	-
Volume - 6th Decile	0.521	(0.118)	-	-
Volume - 7th Decile	0.221	(0.120)	-	-
Volume - 8th Decile	0.191	(0.124)	-	-
Volume - 9th Decile	0.212	(0.117)	-	-
Constant	8.964	(1.093)	9.425	(1.722)
Additional Controls	Forest, Species and Year Dummies		Forest, Species and Year Dummies	
	N = 925		N = 242	
	Wald χ^2 (42)	837.19	Wald χ^2 (33)	178.27
	Prob > chi2	0.000	Prob > chi2	0.000

Note: Table presents the estimates from the binomial entry model, estimated separately on unrestricted large and unrestricted small tracts.

Table 7: 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	2.69	(0.03)	4.93	(0.36)
Avg. Big Bidder Entry	1.55	1.57	(0.00)	0.00	(0.00)
Avg. Total Entry	4.26	4.26	(0.03)	4.93	(0.36)
Avg. Prices	95.12	98.51	(15.51)	95.87	(15.68)
% Sales won by Small Bidders	52.44	54.14	(0.03)	97.93	(0.01)
Avg. Sale Surplus (per mbf)		119.16	(25.85)	100.11	(19.52)
<i>Panel B: Tracts sold by set-aside sale (N=163)</i>					
Avg. Small Bidder Entry	4.50	2.57	(0.05)	4.71	(0.32)
Avg. Big Bidder Entry	0.00	1.71	(0.00)	0.00	(0.00)
Avg. Total Entry	4.50	4.28	(0.05)	4.71	(0.32)
Avg. Prices	90.74	97.29	(17.47)	92.01	(16.87)
% Sales won by Small Bidders	100.00	50.63	(0.04)	98.28	(0.01)
Avg. Sale Surplus (per mbf)		119.21	(29.11)	98.39	(22.00)

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 250 bootstrap repetitions are reported in parentheses.

Table 8: Comparison of Set-Asides and Subsidies

<i>Outcome Variable</i>	Small Bidder Entry	Big Bidder Entry	Total Small Bidder Profit (\$'000s)	Total Big Bidder Profit (\$'000s)	Price (\$/mbf)	Gross Price (\$/mbf)	% Won by Small Bidders	Total Small Bidder Vol.	Surplus (per mbf)
<i>Panel A: All Sales (N=1330)</i>									
Actual Policy Set-Asides Selected by Sale Volume	2.94	1.38	72.78	188.46	97.72	97.72	59.55	2223.64	116.62
Predicted Optimal Set-Asides	3.41	1.04	71.94	189.61	97.46	97.46	68.66	2184.34	113.05
No Subsidy	2.86	1.46	67.39	199.14	98.66	98.66	58.95	2239.06	118.69
5% Subsidy	2.66	1.60	59.24	213.31	98.30	98.30	53.18	1817.95	119.31
6% Subsidy	3.04	1.59	72.84	189.39	101.40	103.07	58.46	2207.21	119.15
10% Subsidy	3.12	1.58	75.30	185.09	101.92	103.95	59.57	2277.82	119.09
15% Subsidy	3.37	1.57	84.15	169.47	103.65	107.34	63.42	2534.93	118.76
20% Subsidy	3.62	1.55	93.46	152.87	105.23	111.29	67.40	2809.66	118.22
	3.82	1.55	101.38	138.60	106.38	115.09	70.81	3043.15	117.58
<i>Panel B: Small Sales (N=265)</i>									
Actual Policy Set-Asides Selected by Sale Volume	3.41	0.27	10.27	0.92	84.66	84.66	88.66	256.54	84.73
Predicted Optimal Set-Asides	3.69	0.00	11.18	0.00	85.30	85.30	96.12	277.92	85.22
No Subsidy	3.65	0.05	11.08	0.13	85.30	85.30	95.09	274.05	85.22
5% Subsidy	3.31	0.33	9.92	1.34	84.42	84.42	85.21	247.26	85.43
6% Subsidy	3.36	0.31	9.97	1.26	84.93	85.57	85.98	250.15	85.48
10% Subsidy	3.41	0.27	10.15	1.08	85.01	85.63	86.94	253.44	85.49
15% Subsidy	3.51	0.17	10.44	0.79	85.25	85.79	89.46	260.97	85.51
20% Subsidy	3.57	0.11	10.62	0.61	85.36	85.71	91.47	265.31	85.46
	3.62	0.07	10.93	0.29	85.46	85.64	93.27	270.16	85.44
<i>Panel C: Large Sales (N=1065)</i>									
Actual Policy Set-Asides Selected by Sale Volume	2.82	1.66	88.33	235.13	100.97	100.97	52.31	2713.11	124.55
Predicted Optimal Set-Asides	3.34	1.30	87.05	236.79	100.49	100.49	61.83	2658.70	119.98
No Subsidy	2.66	1.81	81.40	248.66	101.99	101.99	49.96	2728.01	127.02
5% Subsidy	2.50	1.91	71.51	266.05	101.76	101.76	45.22	2208.78	127.74
6% Subsidy	2.96	1.91	88.48	236.20	105.50	107.42	51.61	2694.17	127.53
10% Subsidy	3.04	1.91	91.51	230.87	106.12	108.51	52.76	2781.54	127.45
15% Subsidy	3.34	1.91	102.49	211.44	108.23	112.70	56.94	3100.75	127.04
20% Subsidy	3.63	1.91	114.07	190.76	110.18	117.65	61.41	3442.76	126.37
	3.87	1.91	123.88	173.02	111.58	122.42	65.22	3733.14	125.58

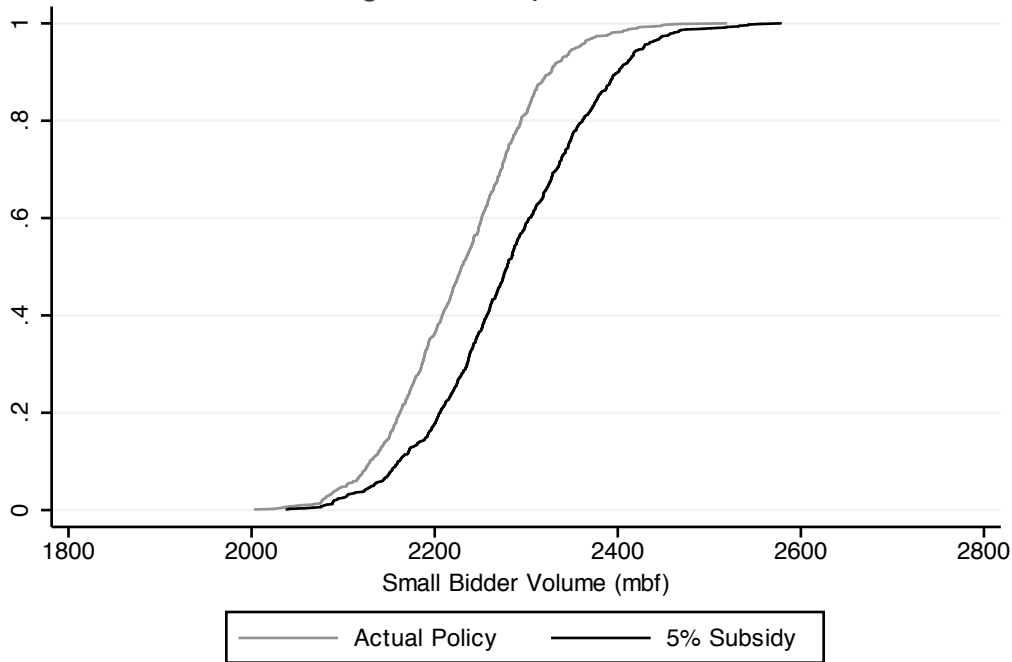
Note: Table reports outcomes, averaged over all auctions, of different set-aside and subsidy policies. The actual outcomes are those given the set-aside policy observed in the data. All set-aside policies shown achieve at least as much total small bidder volume over all sales as actual set-asides do. "Volume-Based 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. "Predicted 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. 469 sales are set aside under volume-based set-asides, and 197 under predicted optimal set-asides. For the counterfactual subsidies, the subsidies are applied to every auction in the sample.

Table 9: 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)	Price (\$/mbf)	Prob. Some Small Bidder Wins	Total Small Bidder Vol.	Surplus (per mbf)
<i>Panel A: Unrestricted Sales with Positive Small Bidder Entry (N=1167)</i>								
No Subsidy or Set-Aside	2.69	1.57	59.13	214.79	98.51	54.14	1802.49	119.16
Set-Aside, Fixed Small Entry	2.69	0.00	176.23	0.00	70.49	87.44	4115.50	86.40
Set-Aside, Endogenous Entry	4.93	0.00	173.23	0.00	95.87	97.93	5020.94	100.11
5 % Subsidy, Fixed Big, Small Entry	2.69	1.57	63.67	207.18	98.93	56.07	1925.40	119.10
5 % Subsidy, Endogenous Entry	3.05	1.58	72.73	190.89	101.53	58.74	2188.01	119.15
<i>Panel B: Set-Aside Sales (N=163)</i>								
No Subsidy or Set-Aside	2.57	1.71	60.47	202.17	97.29	50.63	1939.49	119.21
Set-Aside, Fixed Small Entry	2.57	0.00	178.10	0.00	67.80	88.92	4485.95	85.52
Set-Aside, Endogenous Entry	4.71	0.00	170.49	0.00	92.01	98.28	5238.88	98.39
5 % Subsidy, Fixed Big, Small Entry	2.57	1.71	65.21	194.25	97.75	52.84	2078.83	119.14
5 % Subsidy, Endogenous Entry	2.98	1.70	73.63	178.66	100.51	56.46	2344.64	119.10

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-asides rows, "Fixed Small Entry" refers to fixing small bidder entry probabilities at their equilibrium values in the unrestricted, unsubsidized case, and setting big bidder entry to zero. For the 7.5% subsidy rows, "Fixed Big, Small Entry" refers to fixing big and small bidder entry probabilities 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.

Figure 1: Empirical CDFs



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




Figure 1. Shows the distribution of (per-sale) timber volume won by small bidders across simulations conducted using the actual set-aside policy and a 5% subsidy. To construct the figure, we simulated the whole set of auctions 800 times under each policy. For each simulation, we calculated the total timber volume won by small bidders and divided by the number of sales. The plot shows the distribution of this per-sale volume number.

Figure 2. Shows the allocation from various subsidized and unsubsidized auctions. The x-axis is the highest value among the small bidders, and the y-axis the highest value among the large bidders. The dots on each axis show the deciles of the distributions of these high values. In each case, we assume two big bidders, eight potential small bidders, X equal to the mean among the larger tracts, and $u=1$. The sealed auction cases assume two actual small bidders. The first plot shows the Myerson revenue-optimal allocation with fixed entry, $n=(2,2)$. The final plot shows the revenue-optimal allocation with endogenous equilibrium entry by the bidders.