Quantitative Marketing and Economics (2026) 24:3
https://doi.org/10.1007/s11129-026-09305-1

Affirmative action as a cost cutting tool in procurement markets

Simha Mummalaneni[1]

Received: 3 April 2023 / Accepted: 15 December 2025

Abstract

We study the effect of a procurement market affirmative action program where buyers can provide bid discounts to preferred (small, women-owned, and minority-owned) vendors, thereby enabling them to win contracts even if they are not the cheapest option. These programs are typically framed as an exercise in social responsibility, since buyers will pay slightly more in order to support traditionally disadvantaged businesses. Using data from business-to-government procurement auctions in Virginia and a structural model of vendors' bidding behavior, we demonstrate that preferred vendors generally have higher costs than their non-preferred counterparts. As a consequence, the affirmative action program can reduce buyers' procurement expenditures in this context through intensifying competition and forcing large, low-cost vendors to significantly reduce their prices. The magnitude of these savings is over four times larger when buyers strategically use a variable bid discount policy rather than committing to a fixed bid discount ahead of time. Our findings demonstrate that these affirmative action programs need not be a financial burden for buyers. Instead, affirmative action programs that improve societal outcomes can also improve other key metrics like procurement spending.

Keywords Online platforms · Public policy · Government markets · Diversity · Asymmetric competition · Auctions

JEL codes D44 · H57 · H71 · L38

1 Introduction

Online procurement auctions have become a popular tool for allocating contracts between businesses, as they reduce search frictions and potentially reduce acquisition costs. In the standard setting, the buyer is solely interested in minimizing the costs of

✉ Simha Mummalaneni
nmummalaneni@wm.edu

1 William & Mary, Williamsburg, VA, USA

Published online: 05 March 2026

procurement and therefore awards the contract to the lowest bidder. However, many buyers have additional goals beyond cost minimization, such as supporting local businesses or traditionally disadvantaged vendors as part of an initiative focusing on diversity or other social outcomes. In such instances, buyers modify the traditional procurement auction mechanism by adding affirmative action policies. A common approach is to discount bids from "preferred" vendors; i.e., to treat those bids as being cheaper than they actually are when determining the winner.

Understanding the role of affirmative action in this context is important due to its pervasiveness and the size of the industry: business-to-government procurement accounts for 10 to 20 percent of GDP in countries with developed economies, while business-to-business procurement represents the majority of all economic activity (Dwyer & Tanner, 2002; Hutt & Speh, 2012; OECD, 2014). Affirmative action policies are common in many government markets, ranging from the local to the national level in the United States. However, they are also present in many business-to-business settings. For instance, companies as varied as Chevron, Coca-Cola, Microsoft, and MillerCoors have "Tier II" programs that make it easier for minority vendors to win contracts.

In some of these affirmative action programs, the bid discount level is fixed. For example, the state of California uses a 5 percent bid discount to aid small businesses that bid on government contracts, and South Carolina applies a 7 percent bid discount towards all bids submitted by residents of the state. However, other affirmative action programs use a variable bid discount level — the bid discount is left to the discretion of the buyer and is not announced to vendors at the time of bidding. Vendors in this context know whether or not they are eligible to receive a bid discount, but they do not know the actual bid discount level that is being used.

The goal of this paper is to examine how variable discounting affirmative action programs affect equilibrium outcomes for the buyer. In particular, we focus on how the buyer can set the discount level so as to reduce its overall expenditures while also supporting preferred vendors. By estimating a structural model of vendors' bidding behavior, we can estimate how vendors' bids and buyers' expenditures would change under different bid discount levels. We also examine how these outcomes are affected by the rules determining which types of vendors qualify for the preferred group. These questions are particularly relevant in procurement markets because the buyers typically have the power to define the preferred group and to decide what bid discount level to impose. Therefore, buyers can benefit financially if they better understand the financial impact of the various affirmative action policies that they are considering.

Generally speaking, affirmative action programs are enacted in situations where there are observably different groups and some groups are viewed as deserving preferential treatment. In many contexts, these different groups also have different cost distributions, and the high-cost group is also the preferred group. For example, some affirmative action programs provide bid discounts to small businesses specifically because they tend to have high costs and would find it difficult to win contracts otherwise. In such contexts, the affirmative action program has two opposing potential effects on the buyer's overall procurement expenditures:

1. More auctions are won by preferred (high cost) vendors that did not submit the lowest bid. This leads to an increase in overall expenditures.
2. The affirmative action policy discounts the preferred bids and makes them more competitive, and non-preferred (low cost) vendors respond to the policy by bidding more aggressively than they would have otherwise. This leads to a decrease in overall expenditures.

Our analysis uses bid-level data from the Virginia public procurement market in 2006 and 2007. The buyers in this context are Virginia government agencies who were expected to allocate 40 percent of their procurement dollars to vendors in a preferred group. One major tool at the buyer's disposal was the option to award contracts to vendors from the preferred group who did not submit the lowest overall bid, thereby implying a discretionary level of bid discounting that is not announced to vendors at the time of bidding. Vendors in this context know whether they are eligible to receive a bid discount, but the buyer does not pre-commit to a particular bid discount level, nor does the buyer tell vendors what bid discount level they have chosen. Empirical analysis is especially useful in this situation because asymmetries among bidders make it difficult to predict the direction and magnitude of the affirmative action program's effects.

The auctions in our data are first-price sealed bid auctions, and we observe the full set of bids submitted for each auction. Comparing the cheapest bid to the winning bid in each auction provides an initial examination of how the affirmative action program is implemented. If an auction is won by the lowest bidder, this implies that the bid discount level was too low to play a role in deciding the winner of that auction; this is the case for 81 percent of the auctions in our data. In the remaining 19 percent of auctions that are won by the non-lowest bidder, the cost difference between the winning bid and the cheapest bid is usually less than 8 percent. In about 5 percent of all auctions, however, the winning bid is at least 21 percent more expensive than the cheapest bid. This indicates that there is a fair amount of bid discount heterogeneity across auctions, which is only possible in contexts such as ours where the bid discount level is not fixed.

Since 19 percent of all auctions are won by the non-lowest bidder, an initial interpretation could be that these auctions depict the financial burden that the buyers incur from the affirmative action program — in these auctions, the buyer spent additional money by purchasing from someone other than the cheapest bidder. However, this interpretation does not account for the fact that vendors should respond strategically to changes in the bid discount level by adjusting their bids. Calculating the financial impact of the affirmative action program requires a model that accounts for the fact that the set of bids submitted by vendors is affected by the bid discount level. As a result, we estimate a structural auction model that uses the observed bid data and the equilibrium bidding conditions to estimate the cost distributions for preferred and non-preferred vendors. We then approximate the bidding function for each group at different bid discount levels. Coupling these bidding functions with the estimated costs allows us to simulate the full set of bids and the buyer's overall expenditures under different bid discount levels.

Our dataset includes auctions before and after an important policy change: the preferred group initially consisted of small, women-owned, and minority-owned (SWaM) vendors, but it later changed to only include small businesses. Our structural model allows us to demonstrate that preferred vendors have significantly higher costs than non-preferred vendors in our data context, thereby supporting the intended use of the affirmative action program. We also estimate that the median level of bid discounting is less than one percent under both policy regimes. This level of bid discounting is not sufficient to counteract the asymmetries between the two groups, as non-preferred vendors continue to charge higher markups and enjoy higher profits due to the relatively weak competition.

Finally, we use the estimates from our structural model to examine market outcomes under alternative policy environments. We find that buyers are setting the level of bid discounting too low: higher levels of bid discounting would intensify the level competition, force non-preferred vendors to reduce their markups significantly, and result in the buyer saving money. Our primary contribution is to demonstrate that there can be a significant financial benefit for buyers to use a variable discounting affirmative action policy like Virginia does. In our setting, using a variable discounting policy can lead to a reduction in procurement expenditures that is over four times as large than compared to the best possible fixed discounting alternative.

These financial benefits of the affirmative action are much larger in the first year of our data, when small, women-owned, and minority-owned vendors were included in the preferred group. In that year, the preferred group has much higher costs than the non-preferred group, which means that the affirmative action program can intensify competition and lead to reduced expenditures for the buyer. This pattern of results underlines why a well-designed affirmative action program should encompass two dimensions: the buyer should define the preferred group in a way that results in asymmetric cost distributions between the preferred vs. non-preferred groups, and the buyer should also vary the bid discount level across auctions to yield lower expenditures.

When describing their affirmative action programs, buyers typically frame them as "supplier diversity" initiatives and appeal to broader social responsibility goals. This can lead to tension within the buying organization: social responsibility goals are not always shared by all members of the company, and they can be seen as conflicting with the company's profitability. Various stakeholders may then disagree about whether the merits of the affirmative action program outweigh the added expense. In the case of Virginia's public procurement program, this tension is nicely represented by two quotes from a trade association and a state legislator:

> "Price is important, but not everything. Allow some human equity to be considered when working with diverse vendors."
>
> - Virginia Association of Governmental Purchasing (Fowlkes, 2013)
>
> "A lot of people are concerned about the cost factor. If it's costing the state money, then it's probably worth refining."

- Virginia Delegate Chris Saxman (Kumar, 2008)

Although these two stakeholders disagree about the merits of Virginia's affirmative action program, they are both starting from the shared belief that it will cause the buyer's expenditures to go up — their disagreement comes from whether it is worthwhile to forgo some of their profits in order to increase supplier diversity and support traditionally disadvantaged businesses. However, our findings demonstrate that these programs need not be a financial burden; in fact, a well-designed affirmative action program can lower procurement spending while also allocating more money towards preferred vendors. Therefore, implementing an affirmative action program can allow the buyer to simultaneously both reduce its expenditures and improve other social outcomes of interest.

2 Related literature

In marketing, online procurement auctions have typically been examined in the context of business-to-business markets and industrial procurement. Many of these auctions tend to be "buyer determined" in the sense that the buyer does not commit to focus only on price. Instead, they may prefer a bid that provides the right combinations of price, quality measures, and any other factors that may be relevant. Sometimes the buyer may incorporate these various factors into a transparent scoring rule, and the bid with the best score is awarded the contract. In other situations, the buyer's utility function is private and therefore bidders may not know what the relevant trade-offs are (Haruvy & Jap, 2013; Santamaría, 2015; Stoll & Zöttl, 2017). Another related area of research examines beauty contest auctions in which price is not the sole determinant of the auction allocation and buyers are instead able to use multiple other factors when deciding who wins a particular contract — this is similar to a buyer-determined auction without a scoring rule (Yoganarasimhan, 2013, 2016). Buyer-determined auctions and beauty contest auctions are similar in that they provide buyers with the ability to select a bid that they prefer holistically on multiple dimensions, rather than forcing them to purchase from the cheapest option.

This paper differs from the extant literature on buyer-determined auctions and beauty contest auctions in a few ways. In terms of the marketplace context, one key difference is that buyers in our context are limited by law to only take into account the bid value and the bidder's preferred status. For instance, Yoganarasimhan (2013) finds that non-price factors such as the bidder's physical location and prior interactions with the buyer have important effects on the probability of winning the auction; in our context, these types of factors can be ruled out. Our problem is therefore a more structured one, and accordingly benefits from a different empirical approach. As a consequence, the second difference is that we are examining a distinct question: we show that buyers can use an affirmative action program to intensify competition and reduce their purchasing expenditures, even in situations when they do not (or cannot) place any value on non-price factors. The third difference is that we are able to quantify the additional benefit of a variable bid discounting affirmative action program; i.e., a situation where the buyer can vary the allocation rule across auctions.

This is a departure from the literature on buyer-determined auctions and beauty contest auctions, which typically considers situations where the buyer's utility function (and therefore the buyer's decision making rule) are stable across auctions.

This research uses data from government procurement markets, which represent roughly 10 percent of GDP in the United States, or roughly 2 trillion dollars annually (OECD, 2017). However, there is relatively little marketing research focusing on government procurement markets, despite their substantial financial importance (Grewal & Lilien, 2012; Josephson et al., 2019). More broadly, this research focuses on how buyers can optimally choose the auction rules to yield optimal outcomes, both in terms of their financial expenditures and their social responsibility goals. Therefore, we contribute to the literature studying firms that balance profit seeking with other social goals (Iyer & Soberman, 2016; Sen & Bhattacharya, 2001; Shriver & Srinivasan, 2014), as well as the literature on how to design auction platforms to yield better outcomes for the auctioneer (Cheema et al., 2012; Yao & Mela, 2008).

Economists studying affirmative action programs in an auction setting have typically focused on large-budget projects such as highway and timber auctions (Athey et al., 2013; Krasnokutskaya & Seim, 2011; Marion, 2007; Rosa, 2024). However, there are a few important differences between their papers and this research:

1. They consider settings with a fixed and pre-specified level of bid discounting, as opposed to our setting where the level of bid discounting is neither fixed nor announced to bidders. The variable bid discounting policy we study is common in many private sector business-to-business purchasing contexts, in addition to Virginia's public procurement marketplace. In the counterfactual simulations, we demonstrate that these variable bid discount policies are beneficial for buyers because they can yield significant reductions in procurement expenditures. This is a novel finding that has not been considered by the previous literature.
2. Highway and timber auctions are large-budget items in the hundreds of thousands of dollars. Because of their size and complexity, bidding in highway and timber auctions is complicated by the fact that these jobs are typically parceled out to subcontractors who may have similar agreements with multiple bidders within the same auction. In contrast, we study a setting where most of the auctions are fairly small, with a median unit price of around 20 dollars.
3. Their entry model needs to account for firms making a costly investment to conduct a land survey, discover their own project cost, and come to agreements with subcontractors, especially since the cost of building a highway or logging a timber forest can vary wildly and is dependent on specific local conditions. In our auctions for homogeneous commodity-type products like canned food and office paper, this type of investment is not necessary.
4. Their policy settings are stable, whereas we observe a policy change during the course of our data which alters the composition of the preferred group and indirectly affects the level of bid preference. This exogenous variation allows us to better estimate how changing the composition of the preferred group affects the key outcomes for the buyer, which is an important topic that the previous literature cannot address.

3 Institutional setting

Online procurement auctions have become a popular tool for state governments in recent years, as states have sought to streamline their purchasing processes and cut costs. Vendors benefit from the system because they have a one-stop website that displays all auctions, instead of having to keep track of each buyer's procurement needs separately. Consequently, the buyers are able to receive more bids than before, without having to actively solicit bids from potential vendors. Finally, the public at large benefits because the online auction system helps bring transparency to state spending, thereby reducing the risk of overspending, favoritism, or corruption.

3.1 Virginia's procurement market

Virginia's online procurement system, eVA, was introduced in March 2001. On July 2, 2002, Virginia Governor Mark Warner issued Executive Order 29, which asked the heads of each state agency to provide a written plan explaining how they would "facilitate the participation of small enterprises and enterprises owned by women and minorities in procurement" (Warner, 2002). Soon thereafter, the Virginia government commissioned a report from an independent consulting firm that calculated that just 1.27 percent of state spending was going towards minority- and women-owned businesses (MGT of America Inc., 2004). In response, the Governor's office undertook a series of measures intended to raise the level of state expenditures going towards small, women-owned, and minority-owned (SWaM) vendors. See section A of the online appendix for a more detailed historical summary of the relevant affirmative action policies in Virginia.

For our purposes, the key detail is that by fiscal year 2006, the Governor's office established an aspirational goal of 40 percent of state expenditures going to SWaM vendors and allowed buyers to award contracts to a SWaM vendor even if it was not the overall lowest bidder. However, the policy changed in fiscal year 2007: the "preferred" group now consisted only of small businesses. Woman-owned and minority-owned vendors no longer benefited unless they were also small. The 40 percent aspirational goal remained in place; however, it now meant that buyers were expected to allocate 40 percent of their procurement dollars towards small businesses. See Table 1 for a description of how the preferred group was defined in each year, as well as the size of each group.

Table 1 Mapping of SWaM groups to preferred status, by year

	Preferred	Non-Preferred
2006	Small	
	Women-owned	Non-SWaM
	Minority-owned	
	(n = *206*)	(n = *369*)
2007		Non-SWaM
	Small	Women-owned
		Minority-owned
	(n = *303*)	(n = *522*)

NOTE: The n = _ count denotes the number of vendors in each group in our data

To participate in the state auction system, vendors must register with eVA and submit relevant information, including tax documents. As part of the registration process, vendors specify which product categories they are interested in bidding on.[1] In addition, they can specify whether they are interested in auctions only from specific areas, or whether they are interested in auctions statewide. This filter is especially useful for service-based vendors; for example, a landscaping company may only be interested in contracts that are local. Vendors can later choose to bid on whichever auctions they want to, but submitting this information allows them to receive automatic alerts for auctions that fit the vendor's commodity code availability and its geographic interest. This process reduces the search frictions that would otherwise be present if bidders had to seek these opportunities out one-by-one.

3.2 Data

Our dataset consists of purchases made by Virginia state agencies, public institutions, public universities, and local governments through eVA's Quick Quote system. Quick Quote is an interface that is used for all contracts between \$5,000 and \$50,000 that allows buyers to identify which specific goods they are interested in purchasing. Vendors are competing solely on price, as the state procurement guidelines state that for Quick Quote auctions, "Awards shall only be made on grand total basis" rather than any other holistic criteria (Commonwealth of Virginia: Department of General Services, Division of Purchases and Supply, 2021). Buyers can also choose to use Quick Quote for contracts above \$50,000, but they typically only do so for homogeneous goods. For instance, if a state university is planning on building a new dormitory, they would not use Quick Quote for this purpose, since they would want to make their decision based on a holistic evaluation of each vendor's proposal rather than awarding the project solely on price.

The range of purchases is quite broad. Some of the purchases are common across multiple buyers; for instance, many different buyers are seen purchasing basics like stationery, office equipment, or computers. However, most of the purchases are specific to each buyer's primary mission. Public hospitals purchase medical supplies like syringes, IV tubing, and slings; state parks purchase picnic equipment; and state prisons purchase a wide variety of goods including food, job training books and videos, and sports equipment. Overall, there are 295 different commodity codes that appear in our data. The fifteen most common commodities being purchased are described in Table 2. Many of the auctions are food related, likely reflecting the needs of state-run prisons and universities. However, there is nonetheless a wide variety of commodities appearing in the auction system, even among the most frequent commodities.

Our data consists of Quick Quote auctions made in fiscal years 2006 and 2007, which ran from July 2005 through June 2007. We have information about which buyer was making each request, what items they were requesting, when the bidding period began and ended, how much each participating vendor bid for the contract, and which vendor won the auction. Furthermore, we also have data on the vendor's

[1] These product categories are labeled by the National Institute of Governmental Purchasing (NIGP) commodity codes.

Table 2 Fifteen most frequent commodities

Num. Auctions	Commodity Code	Commodity Description
127	39386	Canned Vegetables (Incl. Canned Salads)
115	38596	Frozen Vegetables
97	38544	Frozen Poultry Entrees
84	20772	Office Printer Accessories and Supplies
74	47017	Canes, Crutches, Gait Trainers, Walkers, etc.
70	87570	Surgical Supplies: Catheters, Needles, Syringes, etc.
65	16500	Commercial Cafeteria And Kitchen Equipment
59	39354	Fruit: Canned, Processed and Preserved
47	39360	Fruit: Juices, Fruit and Vegetable (Not Frozen)
44	38542	Frozen Meat Entrees (Includes Beef and Pork)
41	20186	Female Undergarments and Sleepwear
41	39387	Dried Vegetables: Beans, Peas, etc.
40	39007	Cheese
37	44500	Hand Tools (Powered And Non-Powered)
34	39375	Shortening and Vegetable Oil

Note: The dataset contains 2,331 auctions in total. The fifteen most common commodities listed above account for 975 of those auctions

SWaM status; i.e., whether they were small, woman-owned, minority-owned, or none of the above.

There are 2331 auctions in the data: 755 in 2006 and 1576 in 2007.[2] Table 3 shows that these auctions typically have about five bidders on average: two preferred vendors and three non-preferred vendors. Table 3 also shows that these are typically low price goods, with a median price of less than $20.

Vendors' decisions regarding whether or not to participate could be influenced by the breadth of the affirmative action program; i.e., the definition and size of the preferred category. In our setting, accounting for this is especially important because we observe large shifts in participation after the policy change that narrowed the set of vendors who comprise the preferred category. Table 4 shows the number of bids by group and by year. We see that the number of small bids triples after the policy change, and the number of minority bids declines even despite the overall increase in the number of auctions.

Even though 2006 has a more expansive definition of preferred vendors (comprising all SWaM vendors), 2007 has more preferred bids because many more small bidders entered the market. From the buyer's perspective, a major consequence of the policy change is that it made small vendors more "valuable." In 2006, buyers were

[2] Some Quick Quote auctions are set aside just for small businesses; we drop these set-aside auctions and instead focus only on auctions that were open to all vendors. We do not directly observe why the number of auctions changes so much between years: buyers in 2007 might be purchasing more items, they might be dividing their items across more auctions (i.e., splitting up bulk purchases into multiple smaller purchases), they might be shifting their non-QuickQuote purchases into the QuickQuote system; etc.

Table 3 Summary of auctions: winning bids and number of bids, by year

	Winning bid		Number of bidders	
	Mean	Median	Total	Preferred
2006	\$395	\$16	5.29	2.03
2007	\$442	\$18	4.34	1.92

NOTE: The left half of the table shows the mean and median winning bid across auctions in 2006 and 2007. The right half shows the average number of total bidders and the average number of preferred bidders across auctions in 2006 and 2007

Table 4 Number of bids, by preferred status and year

	SWaM status			Preferred status	
	Small	Woman	Minority	Preferred	Non-preferred
2006	1002	235	297	1534	2462
2007	3032	662	279	3032	3806

NOTE: There are 755 auctions in 2006 and 1576 auctions in 2007. The left half of the table shows the number of bids from small, women-owned, and minority-owned vendors in 2006 and 2007. These categories are not exclusive; a bid from small women-owned vendor would count in both categories. The right half shows the number of bids from preferred and non-preferred vendors in 2006 and 2007

expected to allocate 40 percent of their expenditures towards SWaM vendors, but in 2007, they were expected to allocate 40 percent of their expenditures towards small vendors alone. We expect that buyers would respond by making it easier for preferred bidders to win auctions, and Table 5 indicates that such a shift does in fact occur. In 2006, the preferred and non-preferred bidders have fairly similar probabilities of winning (20% vs. 18%), but these probabilities diverge in 2007 (33% vs. 15%).

4 Model of vendor's bidding

There are two asymmetries between preferred and non-preferred vendors in our context. Cost asymmetries are caused by the inherent differences in vendors' efficiency and scale: large vendors are able to fulfill a given contract more efficiently than a small vendors can. These cost asymmetries are presumably what led Virginia to have such a low percentage of revenue going towards SWaM vendors prior to the enactment of the affirmative action policy, as most small, women-owned, and minority-owned vendors could not adequately compete in a pure lowest-price-wins setting. Payoff asymmetries, on the other hand, are caused by the state's affirmative action

Table 5 Frequency of winning bids, by preferred status and year

	Probability of winning		Number of wins	
	Preferred	Non-preferred	Preferred	Non-preferred
2006	0.20	0.18	303	453
2007	0.33	0.15	988	588

NOTE: The left half of the table shows the probability of winning (the number of auction wins divided by the number of bids) for preferred and non-preferred vendors across auctions in 2006 and 2007. The right half shows the total number of auction wins for preferred and non-preferred vendors across auctions in 2006 and 2007 program. The policy implies that for any given bid level b, a preferred vendor is more likely to win relative to a non-preferred vendor, conditional on them bidding the same amount. In other words:

$$\Pr(b_i \text{ is chosen } | i \text{ is preferred}) \geq \Pr(b_j \text{ is chosen } | j \text{ is non-preferred}) \text{ for } b_i = b_j$$

Therefore, our model of vendor behavior needs to account for two institutional factors: (1) The affirmative action program asymmetrically affects the auction-specific payoffs for each group, and therefore affects the bidding strategy for each group. (2) The cost distributions of preferred and non-preferred vendors are asymmetric.

We assume that for the purposes of evaluating bids, the buyer discounts bids from preferred vendors by dividing the bid by $(1 + \delta)$, where δ is the bid discount value. For instance, if a preferred vendor submits a bit of \$100 and the discount rate is $\delta = 0.10$, then this would be equivalent (from the perspective of the buyer) to a non-preferred vendor submitting a bid of $b = \frac{100}{1.1} = 90.91$. The key assumption here is that a preferred vendor *i* will win a specific auction if two conditions are met: its bid value b_i must be lower than all other preferred bids, and its discounted value must be lower than all non-preferred bids. Formally, b_i wins an auction if and only if:

$$b_i < b_p \text{ for all preferred } p \neq i$$
$$\frac{b_i}{1 + \delta} < b_n \text{ for all non-preferred } n$$

As in Krasnokutskaya and Seim (2011), we are interested in finding group-symmetric equilibria in which all vendors in the same group (preferred or non-preferred) follow the same bidding strategy. This does not mean that they bid the same amount; rather, it means that they use the same mapping of costs to bids – since there are differences in costs, there will therefore be differences in bid values even among vendors from the same preferred category.

Group *n* (non-preferred) does not receive any preferential treatment, while group *p* (preferred) does. We allow for the possibility that cost asymmetries may exist; i.e., that the costs for the two groups are drawn from different distributions with common support. Allowing for asymmetries in our context is important because we observe that bid values tend to vary systematically across groups.[3] Lebrun (2006) proves that there is a unique equilibrium in auctions such as ours where bidders are asymmetric but their cost distributions have common support.

Note that we consolidate our various vendor types (small/woman/minority/none) into two groups: preferred and non-preferred. This step is necessary because vendors in our data do not always have an incentive to get certified in all the categories for which they truly qualify. For instance, a small woman-owned vendor does not benefit from the dual certification – being certified solely as a small business is sufficient to

[3] Formally, we assume that $c_n \sim F_n[c]$, $c_p \sim F_p[c]$, and $c_n, c_p \in [c, \bar{c}]$. Note that this model allows for the possibility that the groups may have asymmetric costs, but it does not impose cost asymmetries. It includes the symmetric costs model ($F_n = F_p$) as a special case.

receive preferred treatment.[4] Focusing simply on preferred and non-preferred vendors allows us to more accurately characterize the bid and cost distributions of those groups, while also addressing the policy questions that are relevant to the buyer.

One important difference between this research and previous studies examining affirmative action programs in procurement auctions (e.g. Athey et al., 2011; Marion, 2007) is that in our context, the bid discount level is not fixed. Instead, buyers are allowed to vary δ across auctions as they see fit. There are a number of plausible factors that could affect the choice of δ:

1. Buyers are expected to spend 40 percent of their money with preferred vendors; higher values of δ will help accomplish this.
2. Buyers have strictly limited budgets, and low values of δ help limit the number of "mis-allocated" contracts that go to the non-lowest bidder.
3. Higher values of δ may encourage non-preferred vendors to bid more aggressively.
4. Lower values of δ may encourage preferred vendors to bid more aggressively.

Our model assumes that buyers choose a bid discount δ after they see the bids in an auction. This potentially allows them to strategically choose auction-specific δ values to improve outcomes of interest: minimizing overall procurement expenditures and meeting the governor's goal of spending 40% of their procurement dollars with preferred vendors. Since the bid discount level for each auction is not announced to vendors, vendors in our model assume that it will be set at the median δ for that year. A key assumption in our model is that all vendors in a particular auction have the same beliefs regarding δ; vendors cannot differ in their belief about how the buyer is going to discount preferred bids. This assumption is required for identification purposes, as we ultimately need to infer the bidders' project costs from their bids. If vendors were allowed to differ in terms of their belief of δ, we would be unable to say whether observed bid differences were due to them having different costs or different beliefs regarding δ.[5] See section B in the online appendix for further discussion and testing regarding this assumption.

We define φ_p, φ_n as the equilibrium inverse bid functions that map bids to costs. For example, for a given preferred vendor i, $\varphi_p(b_i) = c_i$. Denote the number of preferred and non-preferred vendors in a given auction as n_p and n_n, respectively. For a preferred vendor, the profit function is:

[4] See section A of the online appendix for more details about this certification process.

[5] Our model also assumes that vendors' decisions to participate in a given auction does not depend on δ, which implies that entry costs are negligible. We discuss this assumption and provide some validation for it in section B.1 of the online appendix. In addition, we assume that vendors know how many preferred and non-preferred bidders are going to be in each auction, and that vendors know the distributions of costs for preferred and non-preferred vendors (F_p and F_n).

$$
\begin{aligned}
\pi_p(b_i) &= (b_i - c_i)\ \Pr(\text{win} \mid b_i) \\
&= (b_i - c_i)\ \Pr(b_i < \text{ all preferred } b_j)\ \Pr\left(\frac{b_i}{1+\delta} < \text{ all non-pref } b_n\right) \\
&= (b_i - c_i)\left(1 - F_p\left[\varphi_p(b_i)\right]\right)^{n_p - 1}\left(1 - F_n\left[\varphi_n\left(\frac{b_i}{1+\delta}\right)\right]\right)^{n_n}
\end{aligned}
$$

For a non-preferred vendor, the profit function is:

$$
\begin{aligned}
\pi_n(b_i) &= (b_i - c_i)\ \Pr(\text{win} \mid b_i) \\
&= (b_i - c_i)\ \Pr((1+\delta)b_i < \text{ all preferred } b_p)\ \Pr\left(b_i < \text{ all non-pref } b_k\right) \\
&= (b_i - c_i)\left[1 - F_p\left(\varphi_p\left[(1+\delta)b_i\right]\right)\right]^{n_p}\left(1 - F_n\left[\varphi_n\left(b_i\right)\right]\right)^{n_n - 1}
\end{aligned}
$$

The derivative $\frac{\partial \pi_p(b_i)}{\partial b_i}$ yields the first order condition for the preferred group:

$$
1 = \left[b_i - \varphi_p(b_i)\right]\left[\frac{(n_p - 1) f_p\left[\varphi_p(b_i)\right]\varphi_p'(b_i)}{\left(1 - F_p\left[\varphi_p(b_i)\right]\right)} + \frac{n_n f_n\left[\varphi_n\left(\frac{b_i}{1+\delta}\right)\right]\frac{1}{1+\delta}\varphi_n'\left(\frac{b_i}{1+\delta}\right)}{\left(1 - F_n\left[\varphi_n\left(\frac{b_i}{1+\delta}\right)\right]\right)}\right] \quad (1)
$$

Similarly, the derivative $\frac{\partial \pi_n(b_i)}{\partial b_i}$ yields the first order condition for the non-preferred group:

$$
1 = \left[b_i - \varphi_n(b_i)\right]\left[\frac{n_p f_p\left(\varphi_p\left[(1+\delta)b_i\right]\right)(1+\delta)\varphi_p'\left[(1+\delta)b_i\right]}{1 - F_p\left(\varphi_p\left[(1+\delta)b_i\right]\right)} + \frac{(n_n - 1) f_n\left[\varphi_n(b_i)\right]\varphi_n'(b_i)}{\left(1 - F_n\left[\varphi_n(b_i)\right]\right)}\right] \quad (2)
$$

Note that the profit functions and the first order conditions are different for the preferred and the non-preferred groups. Furthermore, these first order conditions do not have a closed form analytic solution and cannot be solved using ordinary differential equation techniques. With symmetric first-price auctions, the bidding function can be reduced to a simple expression consisting of the cost and a markup value; in our context, this is no longer feasible.

5 Estimation

Our estimation can be broken into two separate components:

1. Estimating the distribution of the bid discount level δ_t. This distribution varies by year, and the individual realizations are auction-specific.
2. Estimating a bidding model, thereby allowing us to recover the underlying project costs for the bidders.

It is important to estimate these components in the order above, as the bidding model depends on our estimates of δ_t.

5.1 Distribution of the bid discount level

At the time of bidding, the vendors do not know the specific level of bid discount that the buyer has chosen. However, they know the legal environment: in 2006, they know that small, women-owned, and minority-owned vendors will benefit from bid discounting, and in 2007, they know that only small vendors will benefit. Since they lack full information, vendors must make a boundedly rational guess about δ as they make their bidding decision, and we assume that they use the median value of the δ distribution in that year. Under this assumption, using the observed auction data to infer the yearly distribution of δ also provides us with an estimate of the vendors' yearly beliefs regarding δ.

The parameter δ cannot be negative, as this would imply that non-preferred vendors are the ones receiving the benefits of affirmative action. Therefore, we know that $\delta \in [0, \infty)$. An additional complication is that we do not observe a point estimate for the δ that was used in each auction. Instead, we only observe an interval of values, with three potential options corresponding to three different auction outcomes:

Preferred bid wins with a non-lowest bid: This outcome occurs 19% of the time in our data. If a preferred bid b_p wins, then we know that its discounted bid value must be less than the smallest non-preferred bid b_n:

$$\begin{aligned} \frac{b_p}{1+\delta} &< b_n \\ b_p &< (1+\delta)b_n \\ \frac{b_p}{b_n} - 1 &< \delta \\ \delta &\in \left(\frac{b_p}{b_n} - 1, \infty\right) \end{aligned}$$

This provides us with a lower bound on the auction-specific realization of δ: if the value had been lower than $\frac{b_p}{b_n} - 1$, the winner would have been a non-preferred bidder instead.

Preferred bid wins with the lowest bid: This outcome occurs 37% of the time in our data. If the auction is awarded to a preferred vendor who was the lowest bidder overall, then this particular auction does not tell us anything about δ. The bid discount level could have been anywhere in the interval $\delta \in [0, \infty)$ and this auction outcome would have still occurred.

Non-preferred bid wins: This outcome occurs 44% of the time in our data. If a non-preferred bid b_n wins, then we know that its bid value must be less than the smallest discounted preferred bid b_p:

$$\begin{aligned} b_n &< \frac{b_p}{1+\delta} \\ (1+\delta)b_n &< b_p \\ \delta &< \frac{b_p}{b_n} - 1 \\ \delta &\in \left[0, \frac{b_p}{b_n} - 1\right) \end{aligned}$$

This provides us with an upper bound on the auction-specific realization of δ: if the value had been higher than $\frac{b_p}{b_n} - 1$, the winner would have been a preferred bidder instead.

Out of these three potential outcomes, the latter two represent outcomes where the lowest bidder wins. Only the first outcome (preferred bid wins with a non-lowest bid) represents auctions in which the buyer used its discretionary power to award the contract to a preferred vendor who did not submit the lowest bid, and this happens in only 19% of the auctions in our data.

Table 6 displays some summary stats for the inferred values of $\delta^{\min}$ and $\delta^{\max}$, and Fig. 1 displays trimmed histograms by year. As expected, the $\delta^{\max}$ values are overall larger than the $\delta^{\min}$ values. Both distributions have most of their mass near zero, but they also have a substantial long tail.

For each auction in the data, we use the lowest preferred bid and the lowest non-preferred bid to calculate the interval of δ values that could have rationalized that outcome. If a preferred bid wins without being the lowest bidder, then the interval is $\delta \in \left(\frac{b_p}{b_n} - 1, \infty\right)$. If a preferred bid wins with the lowest bid, then the interval is $\delta \in [0, \infty)$. Finally, if a non-preferred bid wins, then the interval is $\delta \in \left[0, \frac{b_p}{b_n} - 1\right)$.

Once we have the observed auction-specific δ intervals, we can then estimate the overall distribution of δ via maximum likelihood. Denote the endpoints of the observed δ interval as $\delta^{\min}$ and $\delta^{\max}$; for example, if a non-preferred bid wins, then $\delta^{\min} = 0$ and $\delta^{\max} = \frac{b_p}{b_n} - 1$. We model δ as being distributed log-normal with location μ and scale σ, so the likelihood corresponding to a specific auction is

Table 6 Summary statistics for minimum and maximum bid discount levels

	25th pctle	Median	75th pctle
$\delta^{\min}$ (Min. bid discount)	0.027	0.083	0.210
$\delta^{\max}$ (Max. bid discount)	0.044	0.138	0.381

Note: These summary statistics for $\delta^{\min}$ represent auctions in which $\delta^{\min} > 0$, which are auctions where the buyer awarded the auction to a preferred vendor who was not the lowest bidder. The summary statistics for $\delta^{\max}$ represent auctions in which $\delta^{\max} < \infty$, which are auctions where the buyer awarded the auction to a non-preferred vendor

(a) Minimum bid discount level ($\delta^{\min}$) (b) Maximum bid discount level ($\delta^{\max}$)

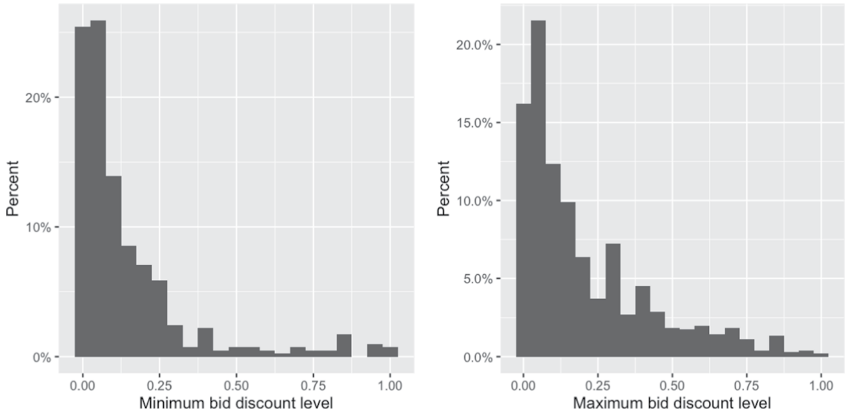

Fig. 1 Trimmed histograms of minimum and maximum bid discount levels. NOTE: Fig. 1a only includes auctions in which $\delta^{\min} > 0$, which are auctions where the buyer awarded the auction to a preferred vendor who was not the lowest bidder. Fig. 1b only includes auctions in which $\delta^{\max} < \infty$, which are auctions where the buyer awarded the auction to a non-preferred vendor

$$\begin{aligned}\ell(\mu, \sigma \mid \delta^{\min}, \delta^{\max}) &= \Pr(\delta^{\min}, \delta^{\max} \mid \mu, \sigma) \\ &= F(\delta^{\max}) - F(\delta^{\min}) \\ &= \Phi\left(\frac{\ln(\delta^{\max}) - \mu}{\sigma}\right) - \Phi\left(\frac{\ln(\delta^{\min}) - \mu}{\sigma}\right)\end{aligned}$$

where $F(\cdot)$ is the log-normal cumulative distribution function and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Note that this approach provides us with a way to estimate the likelihood of observing an interval of values, as opposed to just a single value. Subscripting observations by auction number t allows us to calculate the overall log-likelihood:

$$\mathcal{L} = \sum_{t=1}^{T} \ln\left[\Phi\left(\frac{\ln(\delta_t^{\max}) - \mu}{\sigma}\right) - \Phi\left(\frac{\ln(\delta_t^{\min}) - \mu}{\sigma}\right)\right]$$

We allow μ to vary by year to account for the fact that buyers may adjust their level of bid discounting in response to the policy change. Recall that the policy change shrunk the set of preferred bidders from all SWaM vendors to just small vendors, while at the same time maintaining the expectation that buyers would allocate 40 percent of their money towards preferred vendors. Therefore, we expect that buyers would respond to this by becoming more "aggressive" about giving their money to preferred vendors by increasing their level of bid discounting. See Table 7 for coefficient estimates.

The fact that the μ dummy for 2007 is statistically significant implies that the distribution of δ does in fact vary by year. In fact, it corresponds with our prediction that δ would be higher in 2007 than in 2006, as buyers respond to the policy change. We can see this more clearly in Fig. 2, which displays the estimated density of δ by year.

Table 7 Log-normal estimates for bid discount level

	Estimate		Std. Error
Location μ	−16.34	***	3.82
Location μ dummy for 2007	9.28	***	2.76
Log-Scale $\ln(\sigma)$	2.66	***	0.26

NOTE: Significance levels: *** (0.1 percent), ** (1 percent), and * (5 percent)

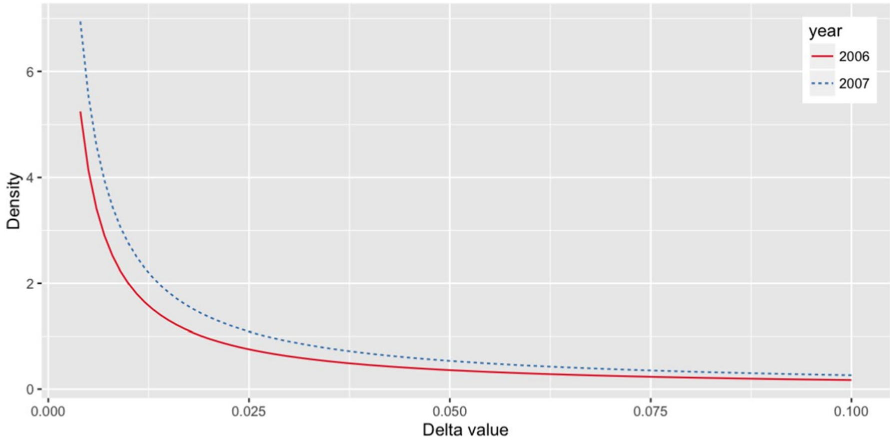

Fig. 2 Density of the bid discount distribution, by year

Figure 2 also allows us to see that the values of δ tend to be quite low. In California, the discount level is set to $\delta = 0.05$ by the government. In our data, the median level of bid discounting per year is:

$$\text{Median}(\delta_{2006}) = \exp(\mu_{2006}) = 8.04 \times 10^{-8} \tag{3}$$

$$\text{Median}(\delta_{2007}) = \exp(\mu_{2007}) = 8.60 \times 10^{-4} \tag{4}$$

Our earlier descriptive statistics indicated that vendors only exercised their discretionary power in 19% of auctions. This pattern is reflected in the estimated δ distributions: the distributions have most of their density at very low values of δ, which implies that in most cases the auction will be awarded to the cheapest bidder unless there is a very small price difference between the cheapest overall bidder and the cheapest preferred bidder.

On the other hand, if a buyer in a particular auction was willing to purchase from a preferred vendor who submitted a much higher bid, then this means that δ for that auction was drawn from the right tail of the overall δ distribution. The significant size of the scale parameter σ means that it is not that surprising to see δ draws above 0.1; according to our estimates, this should happen 16% of the time in 2006 and 37% of the time in 2007. Overall, we can say that the δ distributions in both years have much of their density very close to zero, but that the distributions are wide enough that we occasionally see δ values that are significant enough to make a real difference in terms of deciding the auction winner.

5.2 Distribution of costs

Each observation in the data consists of a bid b_{it} submitted by vendor i in auction t. Each vendor also incurs a cost c_{it} if it wins the auction. Inferring the costs directly would require us to solve the first order conditions for each bidder type (Eqs. 1 and 2). These differential equations do not have a convenient analytic solution in our setting, so we instead adopt the method of Guerre et al. (2000) and rewrite the first order conditions in terms of observables.

There are three steps to the Guerre et al. (2000) procedure: (1) estimate the distribution of bids, (2) use the estimated bid distribution and the inverse bid functions φ_p, φ_n to estimate the costs that can rationalize these two pieces of information, and (3) characterize the distribution of costs. We parameterize the first step with a flexible functional form allowing for unobserved heterogeneity, as in Athey et al. (2011) and Krasnokutskaya and Seim (2011). This allows us to take advantage of the auction characteristics in our data in a parsimonious way.

Define $G_k(b)$ and $g_k(b)$ as the distribution and density of bids for type k. Using the fact that $\varphi_k'(b) = \frac{1}{\left[\varphi_k^{-1}(\varphi_k(b))\right]'}$, these two terms $G_k(b)$ and $g_k(b)$ can be used to replace the unknown model primitives (namely, the distribution of costs) as follows:

$$G_k(b) = F_k\left(\varphi_k(b)\right) \tag{5}$$

$$g_k(b) = \frac{f_k\left(\varphi_k(b)\right)}{\left[\varphi_k^{-1}\left(\varphi_k(b)\right)\right]'} \tag{6}$$

Equation 5 represents the fact that there is a one-to-one relationship between the bid and cost distributions for each group. For example, if vendor A has a higher cost than vendor B and both are in the preferred group, then vendor A will submit a higher bid. Given that our data consists of bids and not costs, the key for us is the reverse inference: if we observe vendor A bidding higher than vendor B, then we also know that vendor A's cost must be higher than vendor B's.

The Guerre et al. (2000) procedure that we adapt does have some limitations. One necessary assumption is that vendors must use the same discount level δ when formulating their bids; otherwise, there would no longer be a one-to-one relationship between the bid and cost distributions, and Eqs. 5 and 6 would no longer hold. In our context, vendors have rational expectations about the median discount level δ; i.e., they use the median δ values calculated in Eqs. 3 and 4 when forming their bids. This assumption means that vendors have the same beliefs about the discount level, and that each vendor's beliefs remains constant across auctions within the same year.[6] Given that vendors have relatively little information about the discount level because

[6] Because vendors use the median bid discount level when forming their bid, the estimation process is not affected by whether buyers are choosing auction-specific bid discounts or whether the bid discounts are fixed throughout the year. As long as the median is the same in both scenarios, the vendors' problem is unchanged and the estimation procedure can recover the correct parameters.

buyers do not directly communicate with them about it, this limitation is realistic in our setting.

We can now find an expression for the inverse bid functions $\varphi_p(b)$ and $\varphi_n(b)$ by re-writing the first order conditions for bidding as :

$$\varphi_p(b_i) : c_i = b_i + \frac{(1+\delta)\left(-1+G_n\left(\frac{b_i}{1+\delta}\right)\right)(-1+G_p(b_i))}{n_n(-1+G_p(b_i))g_n\left(\frac{b_i}{1+\delta}\right)+(1+\delta)(-1+n_p)\left(-1+G_n\left(\frac{b_i}{1+\delta}\right)\right)g_p(b_i)} \tag{7}$$

$$\varphi_n(b_i) : c_i = b_i + \frac{(-1+G_n(b_i))\,(-1+G_p(b_i(1+\delta)))}{(-1+n_n)(-1+G_p(b_i(1+\delta)))g_n(b_i)+(1+\delta)(n_p)\,(-1+G_n(b_i))\,g_p(b_i(1+\delta))} \tag{8}$$

The bid distributions $G_k(b)$ can be estimated directly from the bid data. Combining those bid distributions with the inverse bid functions φ_k (Eqs. 7 and 8) will allow us to infer the cost distributions F_k that rationalize those bids. The cost distributions are the model primitives that we are interested in estimating, because these (when combined with the inverse bid functions φ_k) will allow us to simulate bids and auction outcomes under alternative counterfactual scenarios.

We estimate the bid distributions $G_p(b)$ and $G_n(b)$ parametrically using a Weibull distribution. Our approach allows for unobserved auction-specific characteristics that have an effect on the costs (and therefore the bids) for that auction. This unobserved characteristic has a Gamma distribution and enters multiplicatively into the Weibull bid distribution. Athey et al. (2011) discuss the identification and flexibility of this particular parametric bid distribution, and Krasnokutskaya (2011) demonstrates how to identify the unobserved heterogeneity parameter θ.

The bid distribution includes sets of sale characteristics $\mathbf{X_1}, \mathbf{X_2}$ that are known to both the vendor and the researcher (see Table 8).[7] Any additional heterogeneity across auctions (including across commodity codes) will be captured by the unobservable term *u*.

For a given auction *t*, the distribution of bids for each type *k* (where *k* is either "preferred" or "non-preferred") is:

$$\begin{aligned} G_{k,t}(b) &= 1 - \exp\left[-u_t \cdot \left(\frac{b}{\lambda_k}\right)^{\rho_k}\right] \\ \ln(\lambda_k) &= \beta_0 + \beta\mathbf{X_1} \\ \ln(\rho_k) &= \gamma_0 + \gamma\mathbf{X_2} \\ u_t &\sim \text{Gamma with mean 1 and variance } \theta \end{aligned}$$

The bid parameters β, γ, θ are estimated via maximum likelihood. A particular benefit of this model is that the unobservable term *u* can be integrated out analytically,

[7] We estimated various alternative specifications including models with fixed effects by commodity code, but we found that additional variables had little impact on the overall model fit.

Table 8 Variables entering the bid distribution

$\mathbf{X_1}$ (estimates λ)	$\mathbf{X_2}$ (estimates ρ)
num pref bidders	num pref bidders
num nonpref biddders	num nonpref biddders
dummy (nonpref)	dummy (nonpref)
dummy (2007)	dummy (2007)
dummy (2007 $\times$ pref)	dummy (2007 $\times$ pref)
ln(quantity)	

thereby sparing us from having to numerically integrate over this parameter. As derived in the appendix of Athey et al. (2011), for each auction *t*, the log-likelihood is:

$$\ln(\mathcal{L}_t) = (n_{nt} + n_{pt})\ln(\theta) + \ln\left[\Gamma\left(\frac{1}{\theta} + n_{nt} + n_{pt}\right)\right] - \ln\left[\Gamma\left(\frac{1}{\theta}\right)\right] + \sum_{i=1}^{n_{nt}+n_{pt}} \ln\left[\rho_{it}\frac{1}{\lambda_{it}}\left(\frac{b_{it}}{\lambda_{it}}\right)^{\rho_{it}-1}\right] - \left(\frac{1}{\theta} + n_{nt} + n_{pt}\right)\ln\left[1 + \theta\sum_{i=1}^{n_{nt}+n_{pt}}\left(\frac{b_{it}}{\lambda_{it}}\right)^{\rho_{it}}\right]$$

and the overall log-likelihood is:

$$\ln(\mathcal{L}) = \sum_{t=1}^{T} \ln(\mathcal{L}_t)$$

Estimates of the bid distribution parameters are in Table 9. The fact that our dummy variables are all significant means that (a) bids submitted by preferred bidders and

Table 9 Gamma-Weibull estimates for the bid distribution

	Estimate		Std. Error
	$\ln(\lambda)$		
Constant	3.178	***	0.076
Ln(quantity)	−0.354	***	0.009
Num pref bidders	−0.060	*	0.026
Num non-pref bidders	−0.116	***	0.019
Dummy: non-pref	−0.120	***	0.020
Dummy: year 2007	0.548	***	0.058
Dummy: year 2007 $\times$ pref	−0.109	***	0.022
	$\ln(\rho)$		
Constant	1.355	***	0.026
Num pref bidders	−0.014	*	0.007
Num non-pref bidders	−0.040	***	0.006
Dummy: non-pref	−0.023	***	0.007
Dummy: year 2007	0.250	***	0.018
Dummy: year 2007 $\times$ pref	−0.024	**	0.008
	$\ln(\theta)$		
Constant	2.145	***	0.028

Note: Significance levels: *** (0.1 percent), ** (1 percent), and * (5 percent) non-preferred bidders are significantly different from each other, (b) bid values are significant different across the two years, and (c) the shift in bid values across years is different for each group.

Denote a representative auction as having an average number of bidders, auction characteristics $\mathrm{median}(\mathbf{X_1})$ and $\mathrm{median}(\mathbf{X_2})$, and heterogeneity parameter $u = 1$. Plotting the estimated distributions of bid values $G_k(b)$ for a representative auction allows us to make comparisons across years (see Fig. 3). In both years, non-preferred vendors submit lower bids than preferred vendors. However, there are two major differences across the years. First: the distributions in 2007 are both higher than in 2006, which indicates that after the policy change took place (in particular, after the government narrowed the set of vendors classified as preferred), vendors submitted higher bid values overall. Second: the gap between the groups is much larger in 2006 than in 2007, and this observable difference is confirmed by a Kolmogorov-Smirnov test comparing the two distributions (KS test in 2006: D = 0.074, p-value = 0.0005; KS test in 2007: D = 0.011, p-value = 1).

Although our working hypothesis was that the bid distribution for non-preferred vendors would be lower than the bid distribution for preferred vendors, we did not force this relationship into our estimation. In fact, we did not even force the two groups to have different distributions; we merely included dummy coefficients for groups that could have been zero. The fact that we are able to recover a lower bid distribution for non-preferred bidders compared to preferred bidders is therefore a positive signal for our estimation strategy.

Following Krasnokutskaya (2011), we can recover the cost distribution using the relationship $F_k(c) = G_k(\varphi_k^{-1}(c))$ where the inverse bid function $\varphi_k(\cdot)$ is defined in

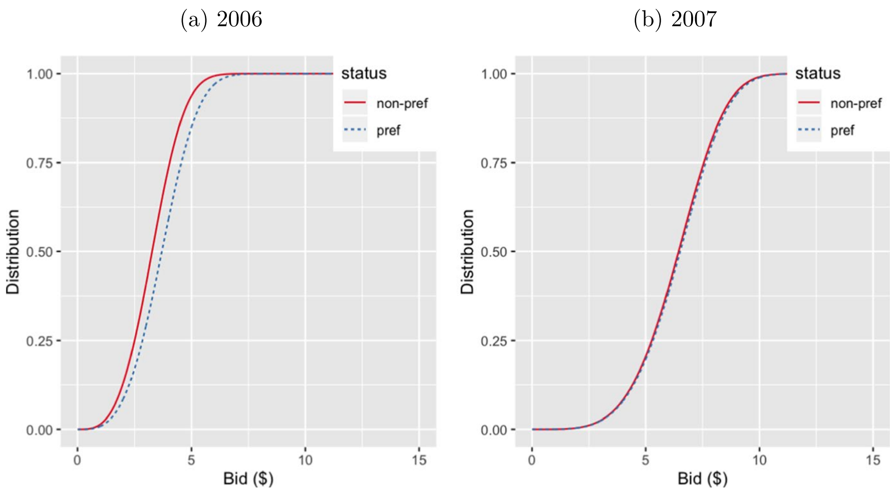

Fig. 3 Estimated bid distribution, by year and group

Eqs. 7 and 8.[8] Figure 4 displays the estimated cost distribution by year and by group. In both years, non-preferred vendors have a lower cost distribution than preferred vendors. As with the estimated bid distributions, we did not impose the condition that non-preferred vendors have lower costs than preferred vendors, so the fact that our estimation procedure recovers this expected relationship is a positive signal. With regards to the policy change, the most notable findings are that the cost distributions for both groups increased after the policy change and that the gap between the two groups' distributions significantly decreased (KS test in 2006: D = 0.263, p-value < 0.0001; KS test in 2007: D = 0.016, p-value = 0.9907). Since we expect that the vendor-specific costs are reasonably stable from year to year, this marked difference is likely due (at least in part) to the policy change.

The policy change had two effects: it changed the composition of each group, and it altered the group-specific incentives to enter. In our context, both of these effects can yield similar patterns: they cause costs to rise for each group. The non-preferred group's costs go up because this group went from consisting of just large (non-SWaM) businesses to now including women- and minority-owned vendors, both of which are likely to have higher costs. Conversely, the preferred group's costs go up because this group went from consisting of small, women-owned, and minority-owned businesses to now only including small businesses. Since small vendors are likely to have higher costs than women- and minority-owned vendors who are not also small, the overall preferred costs go up. Furthermore, the policy change increases the incentive for small vendors to participate, so there are more high-cost small vendors participating in 2007 than in 2006.

This finding holds true even if we assume that each vendor's costs remained the same. In other words, this finding does not mean that there was a cost shock that affected particular vendors or groups. Instead, the shifting of the cost distribution is consistent with the explanation that the policy change altered the definition of the preferred and non-preferred groups and also altered their incentives to enter.

The estimated bid and cost distributions in Figs. 3 and 4 are valid for a representative auction. This representative auction abstracts away from the specific product category, but it does include representative values of the auction characteristics variables $\mathbf{X_1}, \mathbf{X_2}$ listed in Table 8. In addition, the cost distributions depicted in Fig. 4 assume a bid discount δ equal to the median in each year. Our approach of focusing on a representative auction for estimating the structural model and simulating counterfactual outcomes is consistent with Athey et al. (2011). However, as in their paper, our estimation procedure could also be used to recover separate cost distributions for any combination of auction characteristics $\mathbf{X_1}, \mathbf{X_2}$ that a researcher wanted to focus on.

[8] If we did not include the unobserved heterogeneity term u in our estimated bid distribution, we would be able to use the inverse bid functions (Eqs. 7 and 8) to directly infer the individual cost that rationalizes each individual bid.

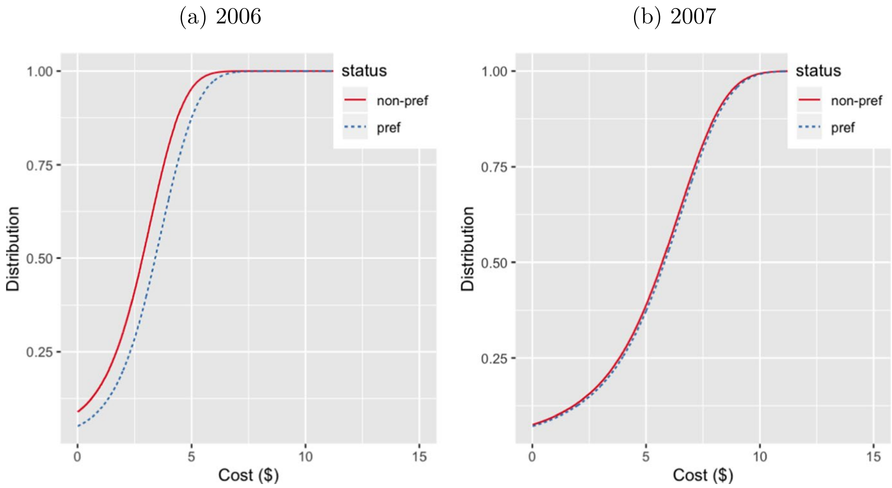

Fig. 4 Estimated cost distribution, by year and group

6 Counterfactual simulations

Having estimated the distribution of costs for preferred and non-preferred vendors, we can now simulate auction outcomes under different policy environments. This is useful because we are interested in examining how vendors and buyers will be affected by different levels of bid discounting. On the vendor's side, we are interested in how vendors of each type strategically respond to the level of bid discounting: to what extent do they increase or decrease their prices as we vary the bid discount? On the buyer's side, we evaluate how the bid discount level affects the total procurement expenditures for the buyer: how much does total spending increase or decrease as we vary the bid discount, and what is the effect of having a fixed vs. variable bid discount?

These questions are particularly relevant in our setting because the level of bid discounting is up to the discretion of each buyer and can vary across auctions. Therefore, it is in the buyer's interest to know how auction outcomes will vary for different bid discount levels. Furthermore, since the state government is balancing the two separate goals of helping preferred vendors win a higher percentage of auction dollars and minimizing overall expenditures, it has a vested interest in finding out how expensive it will be to allocate more business towards these preferred vendors.

6.1 Finding the counterfactual equilibrium

Although the Guerre et al. (2000) approach allows us to estimate the underlying costs for the auction equilibrium, it does not allow us to conduct counterfactual policy simulations. Recall that we were able to estimate the costs by rewriting the bidder's first order condition purely in terms of observables such the bid distribution and the bid density. However, in an alternative policy environment with a different discount level δ, vendors would alter their bidding strategy accordingly. As a result, although the estimated cost distributions are valid, the implied bidding functions in the data would no longer hold.

Therefore, conducting counterfactual simulations requires us to approximate the first order conditions (Eqs. 1 and 2) directly. These first order conditions represent type-symmetric equilibrium bidding behavior. We cannot solve this system using ordinary differential equation techniques; instead, we solve polynomial approximations of the first order conditions (Bajari, 2001). See section C of the online appendix for computational details.

Since we find that the cost distributions are different in 2007 compared to 2006, we estimate two sets of counterfactual simulations: what would happen if we altered δ but kept the 2006 rules in place, and what would happen if we altered δ but kept the 2007 rules in place? This approach is necessary because the effect of the affirmative action policy is highly dependent upon the separation between the two groups' cost distributions: if there is a large gap between observable groups, the affirmative action program can reduce purchasing costs by intensifying the level of competition. However, if the two groups have similar costs, the affirmative action program will be less effective in terms of creating competition, and will instead simply shift contracts to higher-cost vendors.

Given that we have already estimated the costs by group and by year, our procedure is as follows:

1. For a representative auction (as defined in section 5.2) in 2006, draw a cost for each bidder from its 2006 cost distribution F_p or F_n. Mirror this step for a representative auction in 2007.
2. Calculate the optimal bid value for each vendor, based on its own project cost (from step 1), the number of bidders, the distribution of costs, and the bid discount value δ. We estimate bid values for 41 different levels of δ, varying from 0 to 0.40 in increments of 0.01.
3. Find the winner of each auction. The winner is the cheapest overall bid, unless (a) the cheapest bid is from a non-preferred vendor and (b) there is a preferred bid that is within a factor of $(1 + \delta)$ from the cheapest overall bid.
4. Repeat steps 1 - 3 until we have 10,000 auction draws for each year. Since we have 41 potential values of δ and two years in our data, the total number of simulated auctions is $10,000 \times 41 \times 2 = 820,000$.

Although the simulation includes values of $\delta = 0.40$ that are much higher than the typical observed values in our data, this is broadly in line with other government purchasing contexts. For instance, the Department of Defense uses a discount level of $\delta = 0.50$ for domestic vs. foreign vendors (Carpenter & Murrill, 2022).

6.2 Markup values

Since we are interested in seeing how vendors respond to the various bid discount parameter values, we can calculate the average optimal markup percentage for each vendor in each auction, where this metric is defined as $100 \times \frac{b-c}{c}$. Figure 5 displays

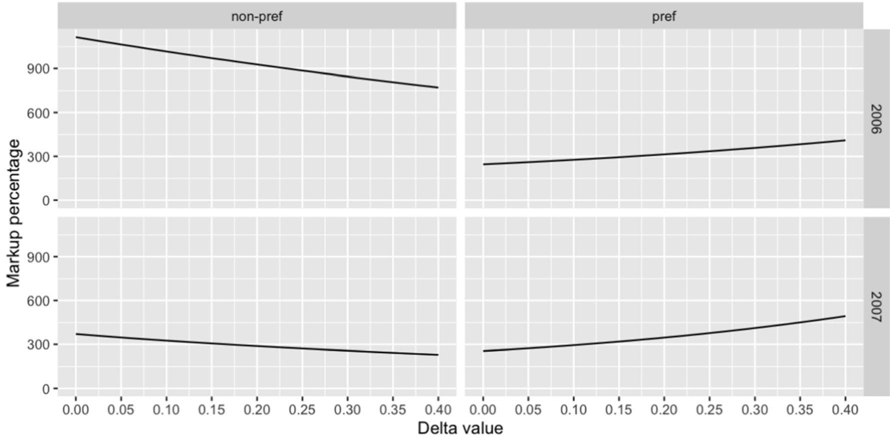

Fig. 5 Average optimal markup percentages, by year and preferred status. NOTE: Markup percentages are defined as $100 \times \frac{b-c}{c}$

a series of plots that show how markup values are affected by δ, for each combination of preferred status and year.

There are notable differences in the optimal markup levels between the two years. In 2006, preferred vendors have lower overall markups than non-preferred vendors, but this pattern is reversed in 2007. These results can be explained by the structural estimates in the previous section. Preferred vendors in 2006 have much higher costs than non-preferred vendors (see Fig. 4a). As a result, we find here that preferred vendors in 2006 are unable to charge high markups, since that would make their bids non-competitive relative to non-preferred vendors. On the other hand, preferred vendors in 2007 have only slightly higher costs than non-preferred vendors (see Fig. 4b). Therefore, they can charge higher markups to take advantage of the affirmative action program, while non-preferred vendors must bid closer to their costs.

The fact that the markup values for preferred and non-preferred vendors in 2007 are nearly equal at $\delta = 0$ serves as a positive sign for the accuracy of our approximation methods. Since the cost distributions for preferred and non-preferred vendors in 2007 are very close to each other, both types of vendors should have similar bidding strategies in a setting without bid discounting. In 2006, however, the two values are further away from each other due to the sizable gap between preferred and non-preferred cost distributions.

More broadly, Fig. 5 demonstrates how the buyer's choice of discount level affects vendors' pricing decisions. In 2006, non-preferred vendors can impose high markups because their significant cost advantages mean that there is limited competition; i.e., they can win even if they do not bid aggressively. This problem is softened as the buyer imposes higher levels of bid discounting: the level of competition intensifies because higher-cost preferred vendors now have a better chance of winning. As a consequence, increases in the discount level force the non-preferred vendors to bid more aggressively (i.e., reduce their prices) to remain competitive.

This result is analogous to prior work on sales contests, in which salespeople compete to out-sell each other for a prize. When there are significant asymmetries in salespeople's ability or the quality of their assigned territories, a traditional sales contest may fail to elicit full effort from its participants – the salespeople who enjoy advantages in ability or territory can win even if they are not exerting full effort (Gopalakrishna et al., 2016; Yang et al., 2013). The firm can mitigate this issue by handicapping the contest (e.g., by providing the weaker salespeople with a head start) so that the stronger salespeople have to expend additional effort in order to win (Ridlon & Shin, 2013; Syam et al., 2013).

In both the sales contest example and our own procurement setting, firms can intervene in an asymmetric market and make it more competitive. Handicapping the strong salespeople or discounting bids from preferred vendors are two ways of providing a more level playing field and increasing the level of competition between salespeople/vendors. However, one key difference is that the affirmative action policies discussed in this paper will occasionally result in contracts being awarded to preferred vendors who are not the cheapest option, thereby incurring a financial loss for the buyer. In the sales contest setting, however, there is no such problem — the firm benefits as long as the contest incentivizes its salespeople to sell more.

6.3 Expenditure minimization

We are interested in discovering which levels of δ yield the lowest procurement cost for the buyer. In other words, how can the government optimally choose its level of bid discounting to reduce its expenditures? The bid discount δ can affect the buyer's expenditures in two ways: it allows preferred vendors to bid less aggressively (to inflate their bid), but it also forces non-preferred vendors to bid more aggressively (to lower their bid). Therefore, the expenditure-minimizing level of δ is that which best balances these two effects.

One major difference between this paper and previous work on affirmative action programs in procurement is that the buyers in our context can vary the bid discounting level from auction to auction. Therefore, we can quantify the financial benefit accorded to Virginia for its variable bid discount policy, relative to a fixed bid discount policy under which buyers would be required to abide by a mandated level of δ. We do so by calculating expenditures under different fixed levels of δ as well as the expenditures if the buyer chooses δ wisely in each auction to minimize the auction-specific expenditures. Under the latter variable discount option, we assume that the buyer picks a median δ value for the year, vendors submit bids for a specific auction, and then the buyer picks the auction-specific δ values that reduce expenditures for that auction while also ensuring that the median δ value for the year matches the number it previously picked.[9] Figure 6 displays results of this exercise.

[9] Our approach for simulating this process works as follows. Imagine that the buyer commits to a median bid discount δ value of 0.05. Out of the 10,000 simulated auction draws, the buyer can abide by this median bid discount by setting $\delta = 0.05$ for 5001 auctions and $\delta = 0$ for the remaining 4999 auctions. For the former set of auctions, the buyer can minimize its expenditures by carefully choosing the 5001 auctions where a δ value of 0.05 has the smallest effect on the final transaction price, because the bid gap between the cheapest non-preferred vendor and the cheapest preferred vendor is low or negative. We implement this rule by calculating the gap between the cheapest non-preferred vendor and the cheapest preferred vendor for each simulated auction, sorting the auctions from lowest to highest on this metric, and

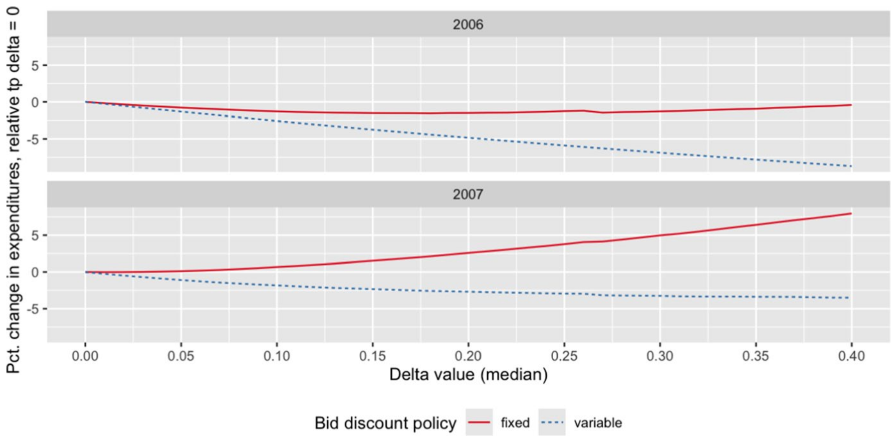

Fig. 6 Change in procurement expenditures under alternative δ values, by year. NOTE: Changes in expenditure values are calculated relative to the $\delta = 0$ condition. The solid red line represents expenditures if the buyer uses an affirmative action policy where the bid discount is fixed, while the dashed blue line corresponds to a policy where the buyer can vary the bid discount across auctions

If the buyer has to abide by a median δ of zero, then there is no ability to intensify competition or induce non-preferred vendors to cut their bids. This is true regardless of whether the buyer uses a fixed or variable bid discount policy, because in our model, vendors are bidding using the median δ value when they form their bids. As a result, we can see in both panels of Fig. 6 that the fixed and variable bid discount policies yield identical expenditures when the median $\delta = 0$.

In Fig. 6, we can see that for all median bid discount values other than $\delta = 0$, the variable bid discount line lies below the fixed bid discount policy line. This means that for any median bid discount value that the buyer wants to abide by, the buyer's costs are lower if they implement a variable bid discount policy rather than a fixed discount. The reason for this is because the variable bid discount policy enables buyers to strategically vary the bid discount. For instance, they can set $\delta = 0$ in situations where a higher bid discount would flip the auction winner from an inexpensive non-preferred vendors to a more expensive preferred vendor. They can also set higher values of δ in situations where doing so would not have this negative financial impact. The key is that vendors are submitting their bids based on the median bid discount for the year but buyers are setting the discount on an auction-by-auction basis.

Although Fig. 6 demonstrates the benefits of a fixed vs. variable bid discount policy from the perspective of the buyer, it does not directly quantify outcomes for different types of vendors. In our context, the primary effect of the fixed vs. variable bid discount policy will be on the percentage of contracts won by preferred vendors.

then using $\delta = 0.05$ for the first 5001 auctions and $\delta = 0$ for the remaining 4999 auctions. This bimodal distribution for δ is a first-best option in terms of allowing the buyer to minimize expenditures while meeting the median δ requirements, but other approaches (such as using a normal distribution for δ) would lead to directionally similar results but with smaller reductions in the expenditures. The key issue is that the more the buyer departs from $\delta = 0$, this increases the number of auctions where the buyer has to pay more than the lowest bid.

We find that this effect is quite small: the percentage of contracts won by preferred vendors declines by about two percentage points (37% to 35%) if the buyer uses a variable vs. fixed bid discount policy in 2006, and there is no effect in 2007. Overall, we conclude that implementing a variable bid discount policy yields slightly worse outcomes for preferred vendors because they win slightly less often, but that these vendor-side outcome differences are much smaller than the ones for the buyer.

Figure 6 shows that the two fixed discount policies can yield small reductions in procurement expenditures relative to the $\delta = 0$ baseline condition. In 2006, the biggest drop in expenditures (1.56%) comes when $\delta = 0.18$, while in 2007, the biggest drop in expenditures (0.016%) comes when $\delta = 0.01$. These expenditure reductions are both much smaller than the comparable numbers in the variable discount scenario: a reduction in procurement expenditures of 8.70% in 2006 and 3.49% in 2007. These results demonstrate that the financial benefits of the affirmative action program are much stronger when buyers can implement a variable bid discount policy rather than a fixed discount policy.

7 Conclusion

This paper examines procurement auction outcomes under a variable discount affirmative action program, in which the buyer does not pre-commit to a particular level of bid discounting. In our context, buyers wield discretion over when and to what extent they discount bids from the preferred category, but they very rarely use this discretionary power. This is likely a result of the buyers' intention to keep procurement costs as low as possible.

We also show that there are significant asymmetries between the cost distributions of preferred and non-preferred vendors in our data context. Therefore, the buyers' reticence to use their affirmative action powers encourages non-preferred vendors to submit bids with high levels of markup. As a result, buyers are being shortsighted with regard to their affirmative action discretion. Increasing their level of bid discounting would reduce overall expenditures by inducing a more robust level of competition among bidders, thereby forcing the low-cost vendors to cut their prices.

Finally, we show that Virginia's variable affirmative action program leads to lower purchasing costs relative to a pre-specified or fixed bid discount. This is due to the fact that buyers can tailor their level of bid discounting given the observed bids that come in. We find that optimally executing this variable bid discount affirmative action program would allow Virginia to lower its procurement costs by about 4 times more relative to the best case scenario in a fixed bid discount environment.

Our structural model is closely tailored to the institution that we study: we allow vendors to have asymmetric costs and payoffs, and we allow for unobservable heterogeneity across auctions. For buyers, we demonstrate that variable bid discount affirmative action programs can reduce overall expenditures, especially when the observable groups have very disparate cost distributions. We also provide evidence that a broad affirmative action program can yield much lower costs of procurement than a narrow one that supports a smaller class of vendors. Finally, from a policy-maker's perspective, we offer evidence that current affirmative action programs are typically not being sold in the best possible light to stockholders, voters, agency heads, and other stakeholders. Typically, these programs are framed as purely an exercise in social responsibility: a way for the buyer to do good by supporting traditionally disadvantaged vendors. However, they should instead be framed as a way to yield social benefit while also lowering expenditures and enabling the buyer to purchase more efficiently.

This research has some limitations that provide opportunities for future research. First, since our data only contains Quick Quote purchases, it is not representative of the full eVA auction system as a whole. Larger projects such as highway procurement and major construction services are not present. This limitation means that we cannot tell whether or when a buyer crosses the 40% threshold in any particular year, and we cannot examine potentially interesting questions like how that threshold affects buyers' allocation decisions or how the buyer's decisions vary as it approaches its yearly budget limit. Second, we focus on a context where buyers are purchasing inexpensive commodity-type products and bidders' entry costs are negligible. Future research could examine the effect of a variable bid discount policy in contexts like highway construction or timber logging where entry costs are substantial; this would require a different model setup and may yield different substantive results.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11129-026-09305-1.

Acknowledgments This research is based on chapter 3 of my dissertation at Northwestern University, and it has benefited greatly from discussions with members of my dissertation committee: Brett Gordon, Song Yao, Rob Porter, Ulf Böckenholt, and Yasutora Watanabe. I thank Ron Bell, Bob Sievert, Shane Caudill, and the staff at eVA and the Virginia Department of General Services for help with accessing the procurement data. I also thank Pradeep Chintagunta, Dominic Coey, Kent Grayson, Ernan Haruvy, Ayelet Israeli, Sandy Jap, Przemek Jeziorski, Jon Levin, Kanishka Misra, Amin Sayedi, Jeff Shulman, Hema Yoganarasimhan, and Jonathan Zhang for their helpful comments and suggestions.

Declarations

Competing Interests The author has no relevant financial or non-financial interests to disclose.

References

Athey, S., Coey, D., & Levin, J. (2013). Set-asides and subsidies in auctions. *American Economic Journal: Microeconomics, 5*(1), 1–27.

Athey, S., Levin, J., & Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics, 126*(1), 207–257.

Bajari, P. (2001). Comparing competition and collusion: A numerical approach. *Economic Theory, 18*, 187–205.

Carpenter, D. H., & Murrill, B. K. (2022). *Congressional research service report r46748: The buy American act and other federal procurement domestic content restrictions*. https://www.congress.gov/crs-product/R46748

Cheema, A., Chakravarti, D., & Sinha, A. R. (2012). Bidding behavior in descending and ascending auctions. *Marketing Science, 31*(5), 779–800.

Commonwealth of Virginia: Department of General Services, Division of Purchases and Supply (2021). *Vendors manual: A vendor's guide on how to do business with the commonwealth of Virginia*. https://dgs.virginia.gov/globalassets/business-units/dps/documents/vendorsmanual/vm-2021.pdf

Dwyer, F. R., & Tanner, J. F. (2002). *Business marketing: Connecting strategy, relationships, and learning*. McGraw-Hill New York.

Fowlkes, L. P. (2013). Shattering 5 myths of diversity programs. *Virginia Association of Governmental Purchasing "News & Views" Newsletter,* 5–6.

Gopalakrishna, S., Garrett, J., Mantrala, M. K., & Sridhar, S. (2016). Assessing sales contest effectiveness: The role of salesperson and sales district characteristics. *Marketing Letters, 27*(3), 589–602.

Grewal, R., & Lilien, G. L. (2012). Business-to-business marketing: Looking back, looking forward. In G. L. Lilien, & R. Grewal (Eds.), *Handbook of business-to-business marketing.* Edward Elgar Publishing.

Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica, 68*(3), 525–574.

Haruvy, E., & Jap, S. (2013). Differentiated bidders and bidding behavior in procurement auctions. *Journal of Marketing Research,50*(2), 241-258.

Hutt, M. D., & Speh, T. W. (2012). *Business marketing management: B2b*. Cengage Learning.

Iyer, G., & Soberman, D. A. (2016). Social responsibility and product innovation. *Marketing Science*.

Josephson, B. W., Lee, J.-Y., Mariadoss, B. J., & Johnson, J. L. (2019). Uncle sam rising: Performance implications of business-to-government relationships. *Journal of Marketing, 83*(1), 51–72.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *The Review of Economic Studies, 78*(1), 293–327.

Krasnokutskaya, E., & Seim, K. (2011). Bid preference programs and participation in highway procurement auctions. *The American Economic Review, 101*(6), 2653–2686.

Kumar, A. (2008). Diversity lags in va. Program to help firms. *The Washington Post*.

Lebrun, B. (2006). Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior, 55*(1), 131–151.

Marion, J. (2007). Are bid preferences benign? The effect of small business subsidies in highway procurement auctions. *Journal of Public Economics, 91*(7), 1591–1624.

MGT of America Inc. (2004). *A procurement disparity study of the commonwealth of Virginia*.

OECD (2014). Size of public procurement market. *Government at a glance 2011*. OECD Publishing. http://dx.doi.org/10.1787/gov_glance-2011-46-en

OECD (2017). Size of public procurement. *Government at a glance 2017*. OECD Publishing. https://doi.org/10.1787/gov_glance-2017-59-en

Ridlon, R., & Shin, J. (2013). Favoring the winner or loser in repeated contests. *Marketing Science, 32*(5), 768–785.

Rosa, B. V. (2024). Subcontracting requirements and the cost of government procurement. *The RAND Journal of Economics, 55*(1), 3–32.

Santamaría, N. (2015). An analysis of scoring and buyer-determined procurement auctions. *Production and Operations Management, 24*(1), 147–158.

Sen, S., & Bhattacharya, C. (2001). Does doing good always lead to doing better? Consumer reactions to corporate social responsibility. *Journal of Marketing Research, 38*(2), 225–243.

Shriver, S. K., & Srinivasan, V. (2014). What if marketers put customers ahead of profits? *Columbia University Working Paper*.

Stoll, S., & Zöttl, G. (2017). Transparency in buyer-determined auctions: Should quality be private or public? *Production and Operations Management, 26*(11), 2006–2032.

Syam, N. B., Hess, J. D., & Yang, Y. (2013). Sales contests versus quotas with imbalanced territories. *Marketing Letters, 24*(3), 229–244.

Warner, M. R. (2002). *Commonwealth of Virginia executive order 29: Equal opportunity in state procurement*. http://digitool1.lva.lib.va.us:1801/webclient/DeliveryManager?pid=42092

Yang, Y., Syam, N. B., & Hess, J. D. (2013). Thrill of victory and agony of defeat: Emotional rewards and sales force compensation. *Quantitative Marketing and Economics, 11*(4), 379–402.

Yao, S., & Mela, C. F. (2008). Online auction demand. *Marketing Science, 27*(5), 861–885.

Yoganarasimhan, H. (2013). The value of reputation in an online freelance marketplace. *Marketing Science, 32*(6), 860–891.

Yoganarasimhan, H. (2016). Estimation of beauty contest auctions. *Marketing Science, 35*(1), 27–54.