

Affirmative Action as a Cost Cutting Tool in Procurement Markets

Simha Mummalaneni *
University of Washington

April 3, 2023

Abstract

We study the effect of a procurement market affirmative action program that grants buyers the ability to purchase from small, women-owned, and minority-owned vendors 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, we find that buyers tend to be unwilling to accept this additional financial burden; instead, they typically only exercise their option to award the contract to the non-cheapest bidder if there is a small gap in bid values. We estimate a structural model of vendors' bidding behavior and demonstrate that there is a significant level of asymmetry between vendors' cost distributions. As a consequence, buyers in this context would reduce their procurement expenditures by roughly 12 percent if they used a stronger affirmative action policy, as this would intensify competition and force large, low-cost vendors to significantly reduce their prices. Therefore, our findings demonstrate that these affirmative action programs need not be a financial burden for buyers. Instead, affirmative action programs that improve diversity and equity outcomes can also improve other key metrics like procurement spending.

Keywords: Online Platforms, Public Policy, Government Markets, Diversity, Asymmetric Competition

*Foster School of Business, University of Washington. Email: simha@uw.edu. 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, Przemysław Jeziorski, Jon Levin, Kanishka Misra, Amin Sayedi, Jeff Shulman, Hema Yoganarasimhan, and Jonathan Zhang for their helpful comments and suggestions.

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 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, equity, and inclusion. 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 and Tanner, 2002; Hutt and 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, most 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 cost asymmetries between groups of bidders and the high-cost group is viewed as deserving preferential treatment. Therefore, the 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 – but not legally obligated – 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. Our focus on affirmative action programs and their implications for vendors and buyers makes this research distinct from extant marketing research on procurement markets (Heide, 2003; Jap and Haruvy, 2008; Noordewier, John, and Nevin, 1990; Silk and Kalwani, 1982).

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 the bid discount levels tend to be low, but that there is

a fair amount of heterogeneity across auctions — such variation is only possible in contexts such as ours where the bid discount level is not fixed.

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. This policy change allows us to better identify how vendors’ decisions to participate and bid in auctions are affected by the competitive environment, and how these vendor decisions affect financial outcomes for the buyer.

The auctions in our data are first-price sealed bid auctions, and we observe the full set of bids submitted for each auction. Since the contracts in our data tend to be for fairly standardized goods and services, our model assumes that vendors know their own cost and also know the distribution of competitors’ costs. Vendors do not know what the particular bid discount level will be for a given auction, but they know based on the policy environment what the median bid discount is for that year.

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 structural estimation approach 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 lack of competition.

Finally, we use the structural estimates from our 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 reduce costs of procurement by nearly 12 percent compared to no affirmative action policy, and by 9 percent even when compared to the best possible fixed discounting alternative. Furthermore, we show that in many cases, the buyer can allocate a significantly higher percentage of expenditures to preferred vendors (e.g., 45 percent instead of 35 percent) without increasing expenditures, thereby implying that many low bid discounting values lead to dominated outcomes.

These substantial benefits of affirmative action are present only 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. When women-owned and minority-owned vendors are removed from the preferred group in the second year of our data, this negates the benefits of the affirmative action program. In that year, the preferred and non-preferred groups have very similar costs, and therefore the buyer cannot create additional competition through the affirmative action program. 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 need not be shared by all members of the company, and they are seen as conflicting with the company’s profitability. Various stakeholders can 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 act in a socially responsible manner.

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 and Jap, 2013; Santamaría, 2015; Stoll and 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 and 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 and Soberman, 2016; Sen and Bhattacharya, 2001; Shriver and Srinivasan, 2014), as well as the literature on how to design auction platforms to yield better outcomes for the auctioneer (Cheema, Chakravarti, and Sinha, 2012; Yao and 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, Coey, and Levin, 2013; Krasnokutskaya and Seim, 2011; Marion, 2007). 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 the norm in many private sector business-to-business purchasing contexts, in addition to Virginia's public procurement marketplace. (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. The benefits of a centralized e-procurement process are fairly clear. 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 woman-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, woman-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 Minority-owned (<i>n</i> = 206)	Non-SWaM (<i>n</i> = 369)
2007	Small (<i>n</i> = 303)	Non-SWaM Women-owned Minority-owned (<i>n</i> = 522)

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

In order to receive and maintain state certification as a small business in Virginia, a vendor either must have 250 or fewer employees or have taken in average revenue of \$10 million or less over the past three years. To qualify as a woman-owned business, a vendor must be at least 51% owned by one or more women and be managed by one or more women. Similarly, to qualify as a minority-owned business, a vendor must be at least 51% owned by one or more racial minorities and be managed by one or more racial minorities. Each vendor can qualify in multiple of these categories if they fulfill the requirements for each. In order to maintain certification for any of these categories, vendors must submit documentation to the state’s Department of Small Business and Vendor Diversity (formerly the Department of Minority Business Enterprise).

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 commodity codes – as defined by the National Institute of Governmental Purchasing (NIGP) – they are able to bid on. 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 on their own.

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 (DGS, 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 will not use Quick Quote for this purpose, since they will be making 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, whether or not the auction was a small business set-aside, how much each participating vendor bid for the contract, and which vendor won the auction. Furthermore, we also have data on the vendor’s SWaM status; i.e., whether they were small, woman-owned, minority-owned, or none of the above. We drop the set-aside auctions and instead focus only on auctions that were open to all vendors.

Since our data only contains Quick Quote purchases, it is not representative of the full

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.

eVA auction system as a whole. Larger projects such as highway procurement and major construction services are not present. This limitation of the data means that we are unable to examine some potentially interesting buyer-side behaviors; for instance, we cannot tell whether or when a buyer crosses the 40% threshold in any particular year, how that threshold affects buyers' allocation decisions, or how the buyer's decisions vary as it approaches its yearly budget limit.

There are 2331 auctions and a total of 10,829 bids in the dataset. As we would expect from an online auction setting, vendors typically bid repeatedly in multiple auctions; there are 1,109 different vendors in our data. Table 3 shows the number of bids per year, broken down by preferred status. There is an increase in bids across both groups which corresponds with the increase in the number of auctions, but the increase is higher (both in absolute and in relative terms) for preferred vendors.

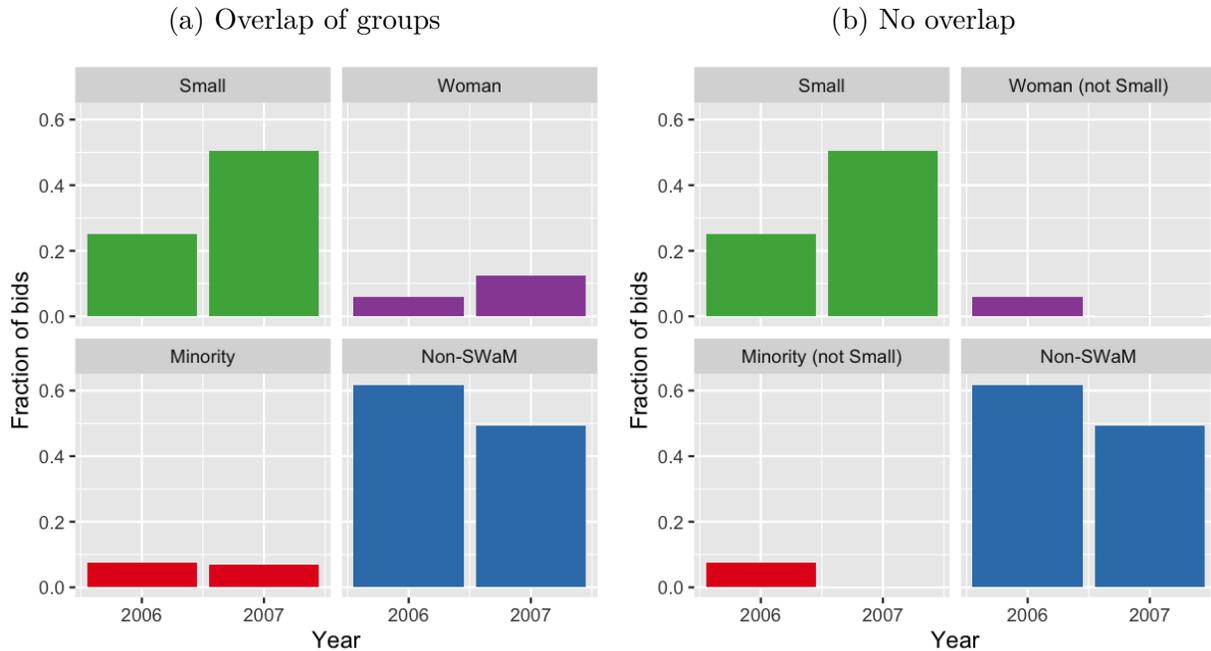
Vendors' decisions regarding whether or not to participate could be influenced by the breadth of the affirmative action program; i.e., the definition and size of the preferred category. In our setting, accounting for this is especially important because we observe large shifts in entry patterns after the policy change that narrowed the set of vendors who comprise the preferred category. Figure 1 shows the fraction of bids by group and by year. One subtlety here is that vendors can be classified into two different SWaM categories if they qualify and register for both. We see that the number of small bids doubles after the policy

Table 3: Number of Bids and Auctions, by year and group

	Number of Bids			Number of Auctions
	Non-Preferred	Preferred	Total	Total
2006	2,460	1,531	3,991	755
2007	3,806	3,032	6,838	1,576

change, and the number of minority and woman bids declines sharply if we consider only vendors who are not also small.

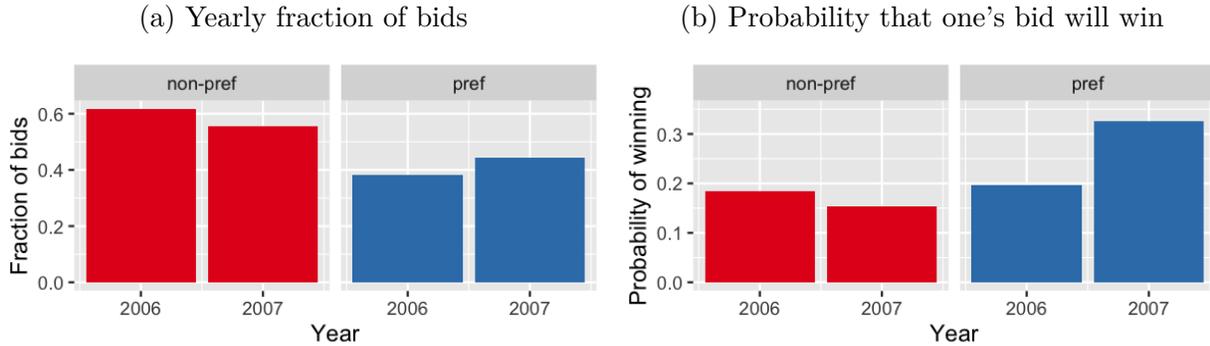
Figure 1: Fraction of bids, by SWaM status and year



NOTE: This figure shows the fraction of bids that are coming from each group. In figure 1a, some vendors are in more than one group; e.g., if they are registered as both small and woman-owned, then they will be counted in both groups. In figure 1b, all small + woman-owned and small + minority-owned vendors are denoted only as small.

Figure 2a displays the fraction of bids that come from preferred vendors and non-preferred vendors on a yearly basis. After the policy change, the prevalence of preferred bids rose about 6 percentage points from 38% to 44%. Comparing this with figure 1 allows us to see that this is due to a large increase in the number of small bids. 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.

Figure 2: Description of bids, by preferred status



From the buyer’s perspective, a major consequence of the policy change is that it made small vendors more “valuable.” In 2006, buyers were 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 small bidders to win auctions, and figure 2b 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 program. The policy implies that for any given bid level b , a preferred vendor bidding b is more likely to win relative to a non-preferred vendor bidding b . 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 bid 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:

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

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.

$$\begin{aligned} c_n &\sim F_n[c] \\ c_p &\sim F_p[c] \\ c_n, c_p &\in [\underline{c}, \bar{c}] \end{aligned}$$

Our model does not impose cost asymmetries – it merely allows for the possibility that the groups may have asymmetric costs. Therefore, this model includes the symmetric costs model ($F_n = F_p$) as a special case. However, allowing for asymmetries in our context is important because we observe that bid values tend to vary systematically across groups. 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 receive preferred treatment. 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, Levin, and Seira, 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.

The fact that δ can vary across auctions means that buyers could potentially adjust δ to account for auction-specific differences in the number of bidders they expect or the relative costs of preferred vs. non-preferred bidders. One key assumption is that buyers choose a value of δ before the auction, but do not announce it. The alternative would be that buyers choose a value of δ upon seeing the bids, which would allow them (for instance) to selectively alter δ to funnel contracts to specific types of vendors or even to specific vendors. This type of behavior would be illegal, which is why we rule it out.

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 . See section B in the online appendix for further discussion and testing regarding this assumption.

We define φ_n, φ_p 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:

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

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

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

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

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

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

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

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.

Our model requires a few commonly-used assumptions regarding vendor behavior:

1. Vendors have identical beliefs about the bid discount level δ .
2. Each vendor's decision to participate in a given auction does not depend on δ .

3. At the time of bidding, each vendor knows how many preferred and non-preferred vendors (n_p and n_n) are going to bid in that auction.
4. Each vendor knows the distributions of costs for preferred and non-preferred vendors (F_p and F_n).

See section B of the online appendix for further discussion of these assumptions.

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. Vendors must make a 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. In other words, we assume that vendors do not know the realized value of δ for a specific auction, but they used their limited information to make a boundedly rational bidding decision. 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. To see this formally, recall that if a preferred bid wins, we know that $\delta \in \left(\frac{b_p}{b_n} - 1, \infty \right)$. If the preferred bid is also the lowest bid overall, then the ratio $\frac{b_p}{b_n}$ is less than one, which means that the lower bound of the interval $\frac{b_p}{b_n} - 1$ is below zero. However, we know that δ cannot be negative, which means that the true interval of plausible values is $\delta \in [0, \infty)$.

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 4 displays some summary stats for the inferred values of δ^{\min} and δ^{\max} , and figure 3 displays trimmed histograms by year. As expected, the δ^{\max} values are overall larger than the δ^{\min} values. Both distributions have most of their mass near zero, but they also have a substantial long tail.

Table 4: Summary Statistics for Minimum and Maximum Bid Discount Levels

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

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

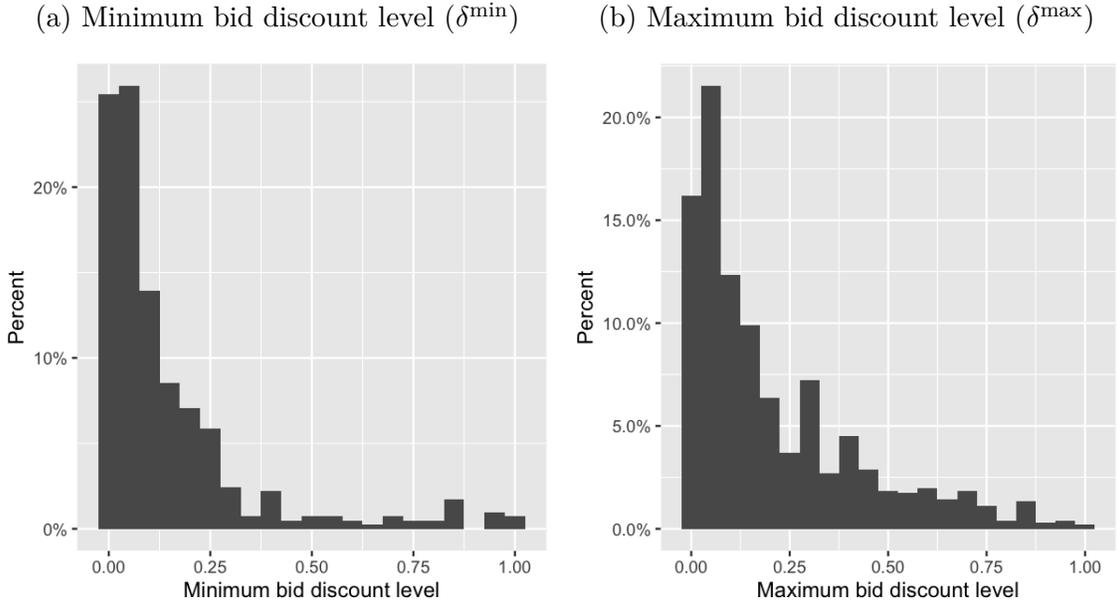
For each auction in the data, we use the lowest preferred bid and the lowest non-preferred bid to calculate the interval of δ values that could have rationalized that outcome. If a preferred bid wins without being the lowest bidder, then the interval is $\delta \in \left(\frac{b_p}{b_n} - 1, \infty\right)$.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

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

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

Figure 3: Trimmed histograms of minimum and maximum bid discount levels



NOTE: Panel 3a only includes auctions in which $\delta^{\min} > 0$; i.e., when the buyer awarded the auction to a preferred vendor who was not the lowest bidder. Panel 3b only includes auctions in which $\delta^{\max} < \infty$; i.e., when the buyer awarded the auction to a non-preferred vendor.

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 5 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 figure 4, which displays the estimated density of δ by year.

Table 5: 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).

Figure 4: Density of the bid discount distribution, by year

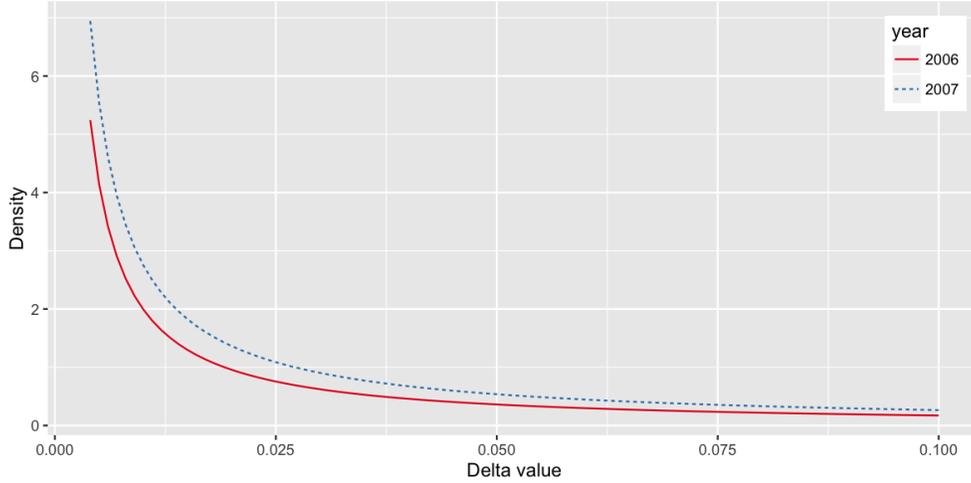


Figure 4 also allows us to see that the values of δ tend to be quite low. In California, the discount level is set to $\delta = 0.05$ by the government. In our data, δ tends to be much lower. The median level of bid discounting per year is:

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

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

These results imply that the buyers are conservative with their use of discretionary affirmative action powers. Generally speaking, they are only willing to award contracts to preferred vendors if their bid is very close to that of the cheapest overall bid. Our earlier descriptive statistics indicated that (a) vendors only exercised their discretionary power in 19% of auctions, and (b) when they do award a contract to the non-lowest bidder, they typically do in cases where the bid is quite close to the cheapest bid. Both of these facts are 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 (equations 1 and 2). These differential equations do not have a convenient analytic solution in our setting, so we instead adopt the method of Guerre, Perrigne, and Vuong (2000) and rewrite the first order conditions in terms of observables.

There are three steps to the Guerre, Perrigne, and Vuong (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, Levin, and Seira (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 $\varphi'_k(b) = \frac{1}{[\varphi_k^{-1}(\varphi_k(b))]'}$, 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:

$$G_k(b) = F_k(\varphi_k(b)) \tag{3}$$

$$g_k(b) = \frac{f_k(\varphi_k(b))}{[\varphi_k^{-1}(\varphi_k(b))]' } \tag{4}$$

Equation 3 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, Perrigne, and Vuong (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 equations 3 and 4 would no longer hold. 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. Given that vendors have relatively little information about the discount level because 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 :

$$\varphi_p(b_i) : c_i = b_i + \frac{(1 + \delta) \left(-1 + G_n\left(\frac{b_i}{1+\delta}\right)\right) (-1 + G_p(b_i))}{n_n(-1 + G_p(b_i))g_n\left(\frac{b_i}{1+\delta}\right) + (1 + \delta)(-1 + n_p) \left(-1 + G_n\left(\frac{b_i}{1+\delta}\right)\right) g_p(b_i)} \quad (5)$$

$$\varphi_n(b_i) : c_i = b_i + \frac{(-1 + G_n(b_i)) (-1 + G_p(b_i(1 + \delta)))}{(-1 + n_n)(-1 + G_p(b_i(1 + \delta)))g_n(b_i) + (1 + \delta)(n_p) (-1 + G_n(b_i)) g_p(b_i(1 + \delta))} \quad (6)$$

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 (equations 5 and 6) 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, Levin, and Seira (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 6). 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. Any additional heterogeneity across auctions

(including across commodity codes) will be captured by the unobservable term u .

Table 6: Variables entering the bid distribution

\mathbf{X}_1 (estimates λ)	\mathbf{X}_2 (estimates ρ)
num pref bidders	num pref bidders
num nonpref bidders	num nonpref bidders
dummy (nonpref)	dummy (nonpref)
dummy (2007)	dummy (2007)
dummy (2007 \times pref)	dummy (2007 \times pref)
$\ln(\text{quantity})$	

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

$$G_{k,t}(b) = 1 - \exp \left[-u_t \cdot \left(\frac{b}{\lambda_k} \right)^{\rho_k} \right]$$

$$\ln(\lambda_k) = \beta_0 + \beta \mathbf{X}_1$$

$$\ln(\rho_k) = \gamma_0 + \gamma \mathbf{X}_2$$

$$u_t \sim \text{Gamma with mean 1 and variance } \theta$$

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

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

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 7. The fact that our dummy variables are all significant means that (a) bids submitted by preferred bidders and 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.

Table 7: 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

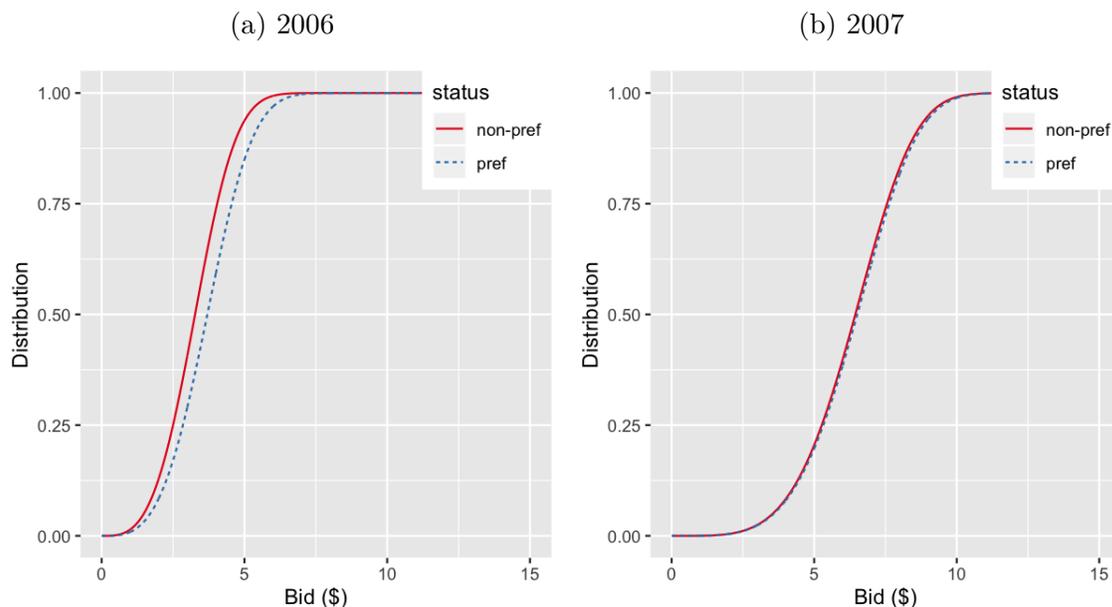
NOTE: Significance levels: *** (0.1 percent), ** (1 percent), and * (5 percent).

Denote a representative auction as having an average number of bidders, auction characteristics $\text{median}(\mathbf{X}_1)$ and $\text{median}(\mathbf{X}_2)$, 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 figure 5). In both years, non-preferred vendors submit lower bids than preferred vendors. However, there are two major differences across the years: the distributions in 2006 are both lower than in 2007, and the gap between the two groups is much larger in 2006 than in 2007. This 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. We try to find the reason for this effect below.

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.

If we did not include the unobserved heterogeneity term u in our estimated bid distri-

Figure 5: Estimated bid distribution, by year and group

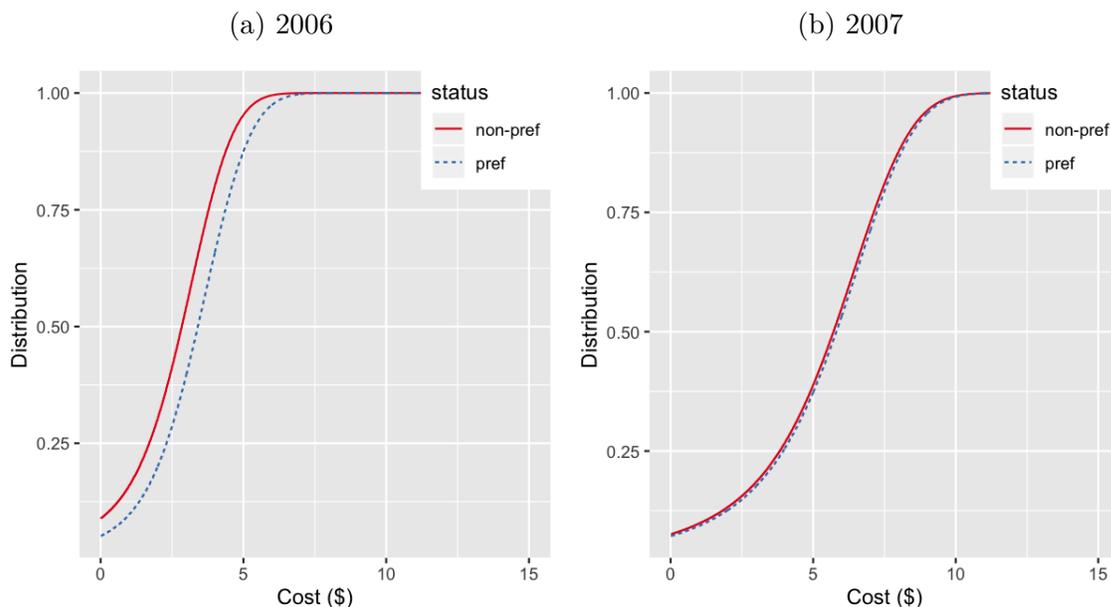


bution, we would be able to use the inverse bid functions (equations 5 and 6) to infer the individual cost that rationalizes each individual bid. Instead, as per Krasnokutskaya (2011), we can recover the cost distribution using the relationship $F_k(c) = G_k(\varphi_k^{-1}(c))$ where the inverse bid function $\varphi_k(\cdot)$ is defined in equations 5 and 6.

Figure 6 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. 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.

Figure 6: Estimated cost distribution, by year and group



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.

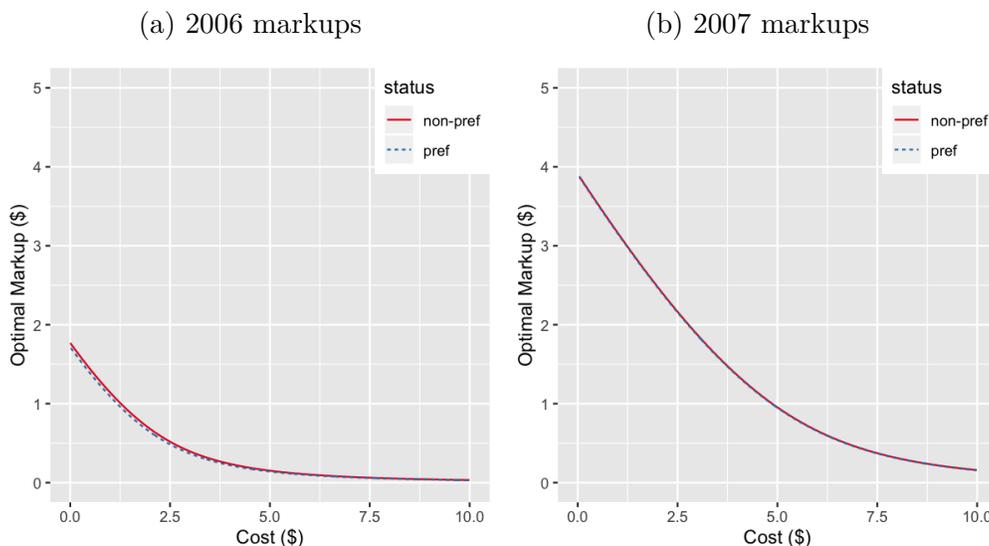
This finding holds true even if we assume that each vendor’s costs remained the same. In other words, this finding does not mean that there was a cost shock that affected particular vendors or groups. Instead, the shifting of the cost distribution is consistent with the explanation that the policy change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

The estimated bid and cost distributions in figures 5 and 6 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 6. In addition, the cost distributions depicted in figure 6 assume a bid discount δ equal to the median in each year. Our approach of focusing on a representative auction for structural estimation and counterfactuals is consistent with Athey, Levin, and Seira (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.

Since we have estimated the distributions of bids and costs, equations 5 and 6 now allow us to estimate the average optimal markup for any given cost. Figure 7 shows how this

markup value varies for each potential cost draw, once again in the context of a representative auction. As we would expect, vendors with low costs can charge high markups and vendors with high costs charge very low markups when submitting their bid. In each year, preferred vendors charge a slightly lower markup value than non-preferred vendors even if they have the same cost – this finding results from the fact that non-preferred vendors with a low cost know that they are unlikely to encounter a preferred vendor with a similarly low cost, and can therefore bid less aggressively. We can see that the gap between each group’s markup values is larger in 2006 than in 2007; this is a result of the fact that there is more of a gap between each group’s cost distribution in 2006 than in 2007 (see figure 6).

Figure 7: Estimated optimal markup values, by year and group



NOTE: Optimal markup for a given cost value is defined as $\text{markup}_k(c) = \varphi_k^{-1}(c) - c$; i.e., the optimal bid value minus the cost. The subscript k denotes the group: preferred or non-preferred.

Figure 7 also implies that the current levels of bid discounting are not sufficient to counteract the natural levels of cost asymmetry between the groups. Despite receiving an advantage from the affirmative action policy, preferred vendors still continue to bid less than non-preferred vendors, even when they have the same cost. This result is because our estimated median bid discount levels are very low. Higher values of δ mean that the payoffs become more asymmetric between groups, which would allow preferred vendors to charge higher markups than non-preferred vendors for any particular cost draw.

Figure 7 demonstrates another important difference across years: for any given cost value, optimal markups are consistently lower in 2006 than in 2007. After the policy change, all vendors (both preferred and non-preferred) choose to bid less aggressively. We posit that this

is due to the entry effects discussed earlier. In 2007, the narrower definition of the preferred group causes more small vendors to register with eVA and fewer women- and minority-owned vendors to register (see figure 1). These small vendors tend to have higher costs than any of the other observable groups in the market, which drives the overall distribution of costs up. As a result, vendors with low cost draws know that they they are in a stronger position and can inflate their bid higher than they otherwise would.

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? Finally, we can measure the financial tradeoff between allocating more business to preferred vendors and trying to minimize overall costs.

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, Perrigne, and Vuong (2000) approach allows us to structurally 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 (equations 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:

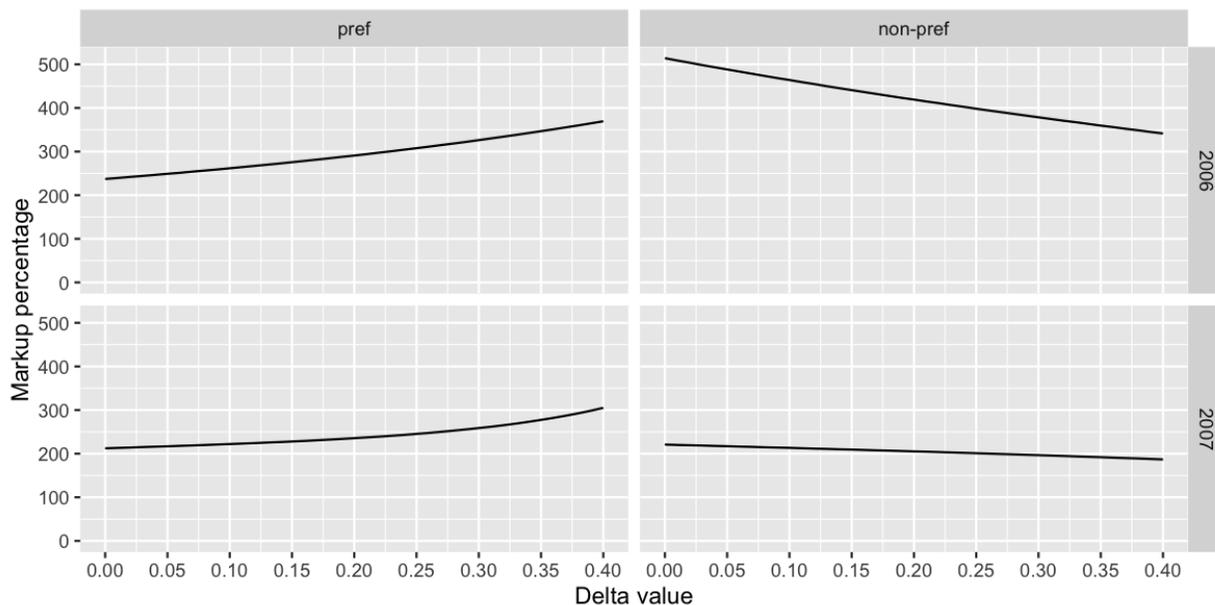
1. For a representative auction (as defined in section 5.2) in 2006, draw a cost for each bidder from its 2006 cost distribution F_p or F_n . Mirror this step for a representative auction in 2007.
2. Calculate the optimal bid value for each vendor, based on its own project cost (from step 1), the number 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.

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

Figure 8: Average optimal markup percentages, by year and preferred status



NOTE: Markup percentages are defined as $100 \times \frac{b-c}{c}$

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 figure 6a). 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 figure 6b). 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, figure 8 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, Syam, and Hess, 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 and Shin, 2013; Syam, Hess, and Yang, 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.

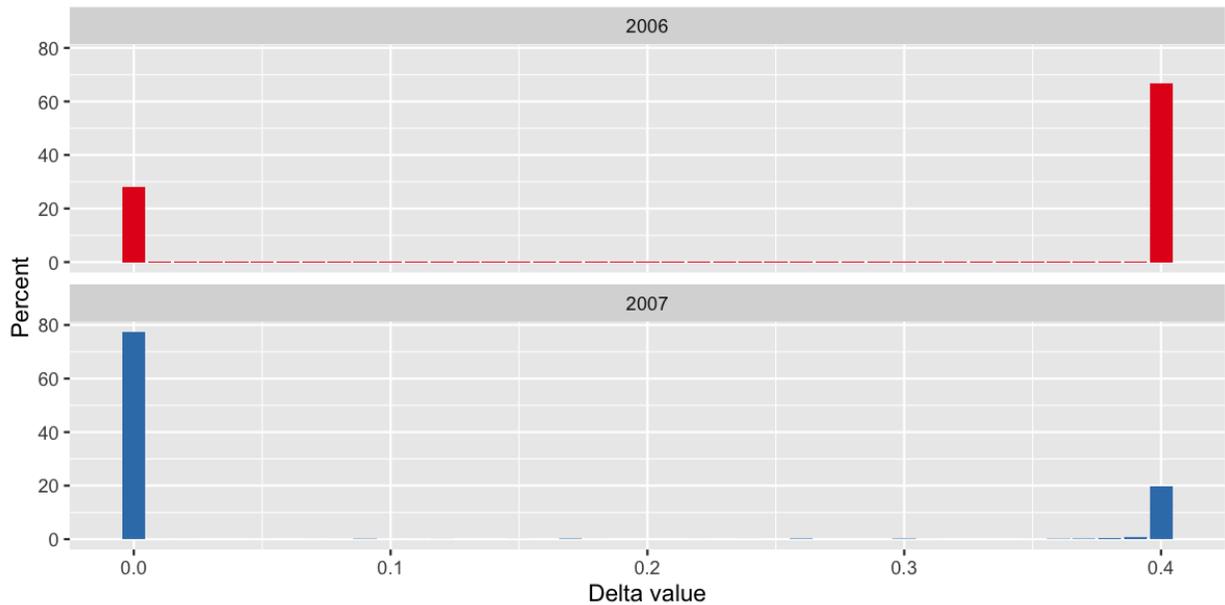
6.3 Expenditure minimization

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 pro-

grams in procurement is that the buyers in our context can vary the bid discounting level from auction to auction. Therefore, we can allow buyers to potentially choose different values of δ for each auction. Figure 9 displays the percentage of times (out of 10,000 simulations) that a given δ value yielded the lowest expenditure.

Figure 9: Frequencies of expenditure-minimizing δ values, by year



NOTE: Values represent the percentage of auction draws in which a given δ value yielded the lowest expenditure values for that auction.

The results are striking for two reasons: the distribution of optimal δ values varies significantly across years, and both distributions are noticeably different from the δ distribution that we estimated using the data (see figure 4). The fact that they differ so strongly from the estimated distribution supports our earlier claim that buyers are under-utilizing their discretionary power – even if they do not care about helping preferred vendors, engaging in higher levels of bid discounting would be in their own financial self interest. This finding is especially stark in 2006, where we find that the buyer can minimize expenditures if it sets $\delta = 0.40$ in the majority of cases.

Figure 9 demonstrates the potential benefits of Virginia’s variable affirmative action program. In some auctions, the vendors’ costs will be such that there is no financial benefit from having a bid discount: perhaps the lowest-cost vendor is a small vendor, or alternatively there might be a negligible gap in costs between the lowest-cost vendor and the lowest-cost preferred vendor. In both of these instances, the affirmative action program will not yield financial benefits for the buyer, and the buyer would prefer to set $\delta = 0$. On the other hand, in auctions where there is a sizable cost gap between the cheapest bid and the cheapest

preferred bid, the buyer can benefit from setting a higher bid discount value. However, we expect that the buyer is unlikely to truly have full information about the vendors' costs, so this analysis is a "first-best" hypothetical scenario.

The difference in distributions of expenditure-minimizing δ across the two years is attributable to the fact that the cost distributions have changed. The affirmative action policy can lower costs if there is a significant gap between preferred and non-preferred cost distributions – this is certainly true in 2006, but less so in 2007 (see figure 6). Even in 2007, though, the optimal distribution of δ is higher than the estimated distribution of δ .

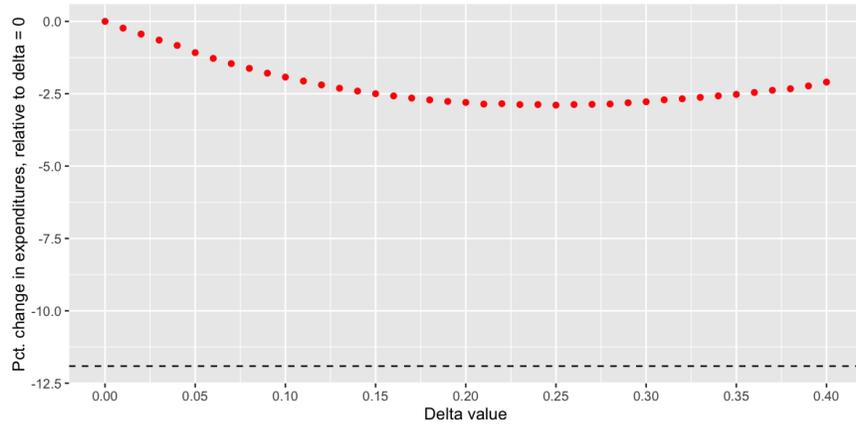
We see from figure 9 that the optimal (expenditure-minimizing) level of δ varies across auctions, even within the same policy regime. Therefore, we can quantify the financial benefit accorded to Virginia for its variable bid discount policy, relative to a fixed bid discount policy under which buyers 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. This can be interpreted as the first-best cost minimization policy for the buyer, since it requires the buyer to have full information about the costs of participating vendors. Figure 10 displays results of this exercise.

The shape of the two curves in figure 10 differ in important ways. In 2006, the expenditures have a U-shaped pattern with respect to δ , while in 2007, they consistently increase. This difference is due to the cost structure in each year: if the preferred group has much higher costs than the non-preferred groups (as is the case in 2006; see figure 6), then higher values of δ can reduce procurement costs by forcing strong bidders to bid more aggressively. Conversely, if the two groups have similar costs (as is the case in 2007), then higher values of δ will increase procurement costs because the level of competition cannot be intensified further. We see these patterns reflected in our results, and we find that the financial effect of the affirmative action program is quite different across years. In 2006, choosing $\delta = 0.25$ across the board reduces expenditures by about 3 percent relative to choosing $\delta = 0$. In 2007, the overall expenditures are lowest at $\delta = 0$.

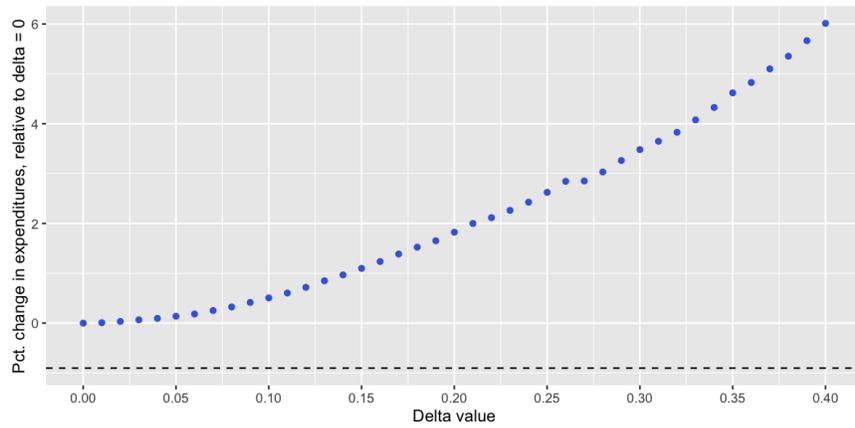
Figure 10 also shows that the "optimal" auction-varying δ rule leads to noticeably lower expenditures when compared to any of the fixed-discount alternatives. With these estimates, we can now calculate the value of adopting a variable bid discounting affirmative action program. We make two comparisons: to the case where there is no affirmative action policy (i.e., $\delta = 0$), and to the case where the buyer uses a fixed discount affirmative action policy. In the latter case, we use a very conservative approach by assuming that the buyer would otherwise choose δ that lowered across-the-board expenditures, despite the fact that we observe the buyers in our data setting δ too low. The cost reductions are substantial in

Figure 10: Change in expenditures under alternate δ values, by year

(a) 2006



(b) 2007



NOTE: Expenditure values are calculated relative to the $\delta = 0$ condition. The dashed lines correspond to the first-best variable discount outcome in which the buyer optimally varies δ by auction to minimize expenditures.

2006, but fairly minimal in 2007.

Cost decrease in 2006 (relative to no affirmative action policy) = 11.91%

Cost decrease in 2006 (relative to fixed affirmative action policy) = 9.29%

Cost decrease in 2007 (relative to no affirmative action policy) = 0.90%

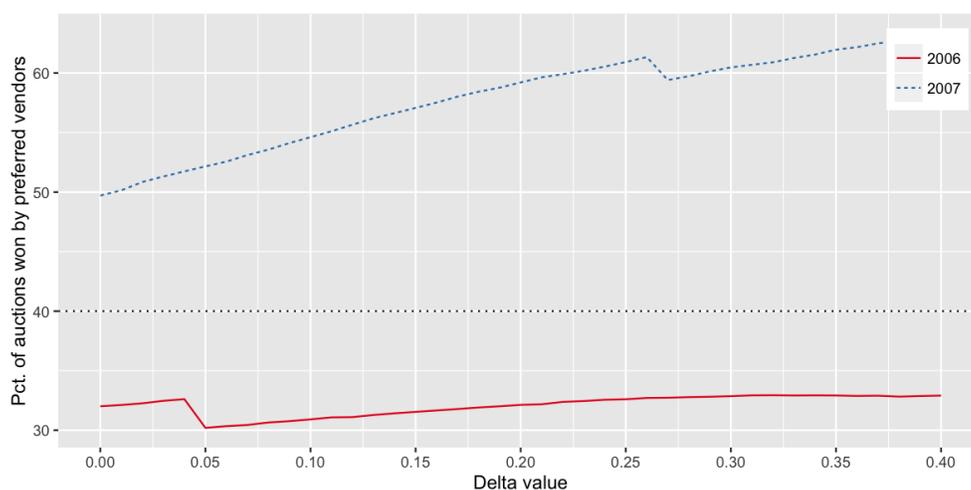
Cost decrease in 2007 (relative to fixed affirmative action policy) = 0.90%

Based on these findings, we conclude that the variable bid discount policy (relative to a fixed policy) yields much larger financial benefits in 2006 than in 2007. This result is driven by the differences in cost structures across the two years, and it underlines the benefit (to the buyer) of having asymmetric costs between preferred and non-preferred vendors. The gap in costs between preferred and non-preferred vendors is much larger in 2006 than in 2007. Therefore, in 2006, the buyer can lower its expenditures by setting high values of δ in many auctions. This is generally not true in 2007, which is why there are minimal gains to having a variable bid discounting level in that year.

6.4 Preferred vendors' revenues

The affirmative action program in Virginia was intended to help preferred vendors who would otherwise find it difficult to compete against their non-preferred counterparts. We can now evaluate how varying the level of bid discounting affects auction outcomes for these preferred vendors. Figure 11 displays the probability that an auction is won by a preferred vendor, depending on the value of δ .

Figure 11: Probability that a preferred vendor wins under alternate δ values, by year

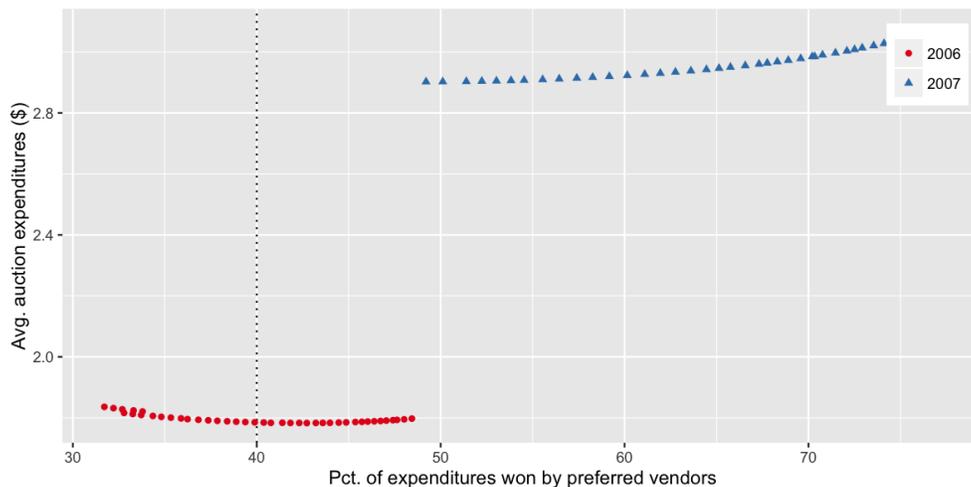


NOTE: The dotted line corresponds to 40% of auctions being won by preferred vendors

The difference between the two years is significant: preferred vendors win a much higher percentage of auctions in 2007 than in 2006, regardless of the bid discount level. The win probabilities are affected by two different forces. At very low levels of δ , preferred vendors are relatively unlikely to win auctions because their costs are higher than those of non-preferred vendors and they are not receiving any special treatment. At very high levels of δ , preferred vendors' probability of winning is moderated by the fact that the level of bid discounting is so substantial that their profit-maximizing incentive is to charge a very high markup – this means that they win fewer auctions, but receive high profits when they do win.

Buyers in this setting have two goals when they choose the bid discount δ : to keep their expenditures low, and to allocate 40% of their overall procurement expenditures to preferred vendors. Figure 12 demonstrates how these two outcomes are related to each other.

Figure 12: Total expenditures vs. pct. of expenditures won by preferred vendors, by year



NOTE: The dotted line corresponds to 40% of expenditures being won by preferred vendors. Expenditures are based on 10,000 auction draws for each δ value

Note that the buyer's underlying choice variable δ is driving both the x- and y- variables in figure 12. This figure is notable because it demonstrates that there are a number of win-win outcomes for the buyer where it can simultaneously lower its expenditures while also allocating more money towards preferred vendors. For instance, we can see that in 2006, the buyer's expenditures are lower if it chooses δ to allocate 40% of expenditures towards small businesses instead of choosing δ to allocate 35% of expenditures towards small businesses. Although we do not impose a specific utility function for the buyer, the former option dominates the latter as long as the buyer places some non-negative value on both cost minimization and helping preferred vendors. The existence of these win-win outcomes is the clearest sign that buyers are not currently fully appreciating the potential benefits of

utilizing the affirmative action program.

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 nearly 12 percent 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 if they have some knowledge about how many bidders are likely to enter and what their cost distributions look like. We find that optimally executing this variable bid discount affirmative action program would allow Virginia to lower its procurement costs by about 9 percent relative to the best case scenario in a fixed bid discount environment.

Our structural approach 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.

Funding and Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

References

- Athey, S., D. Coey, and J. Levin (2013). “Set-Asides and Subsidies in Auctions”. *American Economic Journal: Microeconomics* 5 (1), pp. 1–27.
- Athey, S., J. Levin, and E. Seira (2011). “Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions”. *Quarterly Journal of Economics* 126 (1), pp. 207–257.
- Bajari, P. (2001). “Comparing Competition and Collusion: A Numerical Approach”. *Economic Theory* 18, pp. 187–205.
- Cheema, A., D. Chakravarti, and A. R. Sinha (2012). “Bidding Behavior in Descending and Ascending Auctions”. *Marketing Science* 31 (5), pp. 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. and J. F. Tanner (2002). *Business marketing: Connecting strategy, relationships, and learning*. McGraw-Hill New York.
- Fowlkes, L. P. (June 2013). “Shattering 5 Myths of Diversity Programs”. *Virginia Association of Governmental Purchasing “News & Views” Newsletter*. Ed. by L. B. Smith, pp. 5–6.
- Gopalakrishna, S., J. Garrett, M. K. Mantrala, and S. Sridhar (2016). “Assessing sales contest effectiveness: the role of salesperson and sales district characteristics”. *Marketing Letters* 27 (3), pp. 589–602.
- Grewal, R. and G. L. Lilien (2012). “Business-to-Business Marketing: Looking Back, Looking Forward”. In: *Handbook of Business-to-Business Marketing*. Ed. by G. L. Lilien and R. Grewal. Northampton MA: Edward Elgar Publishing.
- Guerre, E., I. Perrigne, and Q. Vuong (2000). “Optimal Nonparametric Estimation of First-Price Auctions”. *Econometrica* 68 (3), pp. 525–574.
- Haruvy, E. and S. Jap (2013). “Differentiated Bidders and Bidding Behavior in Procurement Auctions”. *Journal of Marketing Research* 50 (2), pp. 241–258.
- Heide, J. B. (2003). “Plural Governance in Industrial Purchasing”. *Journal of Marketing* 67 (4), pp. 18–29.
- Hutt, M. D. and T. W. Speh (2012). *Business marketing management: B2B*. Cengage Learning.

- Iyer, G. and D. A. Soberman (2016). “Social Responsibility and Product Innovation”. *Marketing Science*.
- Jap, S. and E. Haruvy (2008). “Interorganizational Relationships and Bidding Behavior in Industrial Online Reverse Auctions”. *Journal of Marketing Research* 45 (5), pp. 550–561.
- Josephson, B. W., J.-Y. Lee, B. J. Mariadoss, and J. L. Johnson (2019). “Uncle Sam Rising: Performance Implications of Business-to-Government Relationships”. *Journal of Marketing* 83 (1), pp. 51–72.
- Krasnokutskaya, E. (2011). “Identification and Estimation of Auction Models with Unobserved Heterogeneity”. *The Review of Economic Studies* 78 (1), pp. 293–327.
- Krasnokutskaya, E. and K. Seim (2011). “Bid Preference Programs and Participation in Highway Procurement Auctions”. *The American Economic Review* 101 (6), pp. 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), pp. 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), pp. 1591–1624.
- MGT of America Inc. (2004). *A Procurement Disparity Study of the Commonwealth of Virginia*.
- Noordewier, T. G., G. John, and J. R. Nevin (1990). “Performance Outcomes of Purchasing Arrangements in Industrial Buyer-Vendor Relationships”. *Journal of Marketing*, pp. 80–93.
- OECD (2014). “Size of Public Procurement Market”. In: *Government at a Glance 2011*. http://dx.doi.org/10.1787/gov_glance-2011-46-en. OECD Publishing.
- (2017). “Size of Public Procurement”. In: *Government at a Glance 2017*. https://doi.org/10.1787/gov_glance-2017-59-en. OECD Publishing.
- Ridlon, R. and J. Shin (2013). “Favoring the winner or loser in repeated contests”. *Marketing Science* 32 (5), pp. 768–785.
- Santamaría, N. (2015). “An analysis of scoring and buyer-determined procurement auctions”. *Production and Operations Management* 24 (1), pp. 147–158.
- Sen, S. and C. Bhattacharya (2001). “Does Doing Good Always Lead to Doing Better? Consumer Reactions to Corporate Social Responsibility”. *Journal of Marketing Research* 38 (2), pp. 225–243.
- Shriver, S. K. and V. Srinivasan (2014). “What If Marketers Put Customers Ahead of Profits?” *Columbia University Working Paper*.

- Silk, A. J. and M. U. Kalwani (1982). “Measuring Influence in Organizational Purchase Decisions”. *Journal of Marketing Research*, pp. 165–181.
- Stoll, S. and G. Zöttl (2017). “Transparency in Buyer-Determined Auctions: Should Quality be Private or Public?” *Production and Operations Management* 26 (11), pp. 2006–2032.
- Syam, N. B., J. D. Hess, and Y. Yang (2013). “Sales contests versus quotas with imbalanced territories”. *Marketing Letters* 24 (3), pp. 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., N. B. Syam, and J. D. Hess (2013). “Thrill of victory and agony of defeat: Emotional rewards and sales force compensation”. *Quantitative Marketing and Economics* 11 (4), pp. 379–402.
- Yao, S. and C. F. Mela (2008). “Online Auction Demand”. *Marketing Science* 27 (5), pp. 861–885.
- Yoganarasimhan, H. (2013). “The Value of Reputation in an Online Freelance Marketplace”. *Marketing Science* 32 (6), pp. 860–891.
- (2016). “Estimation of Beauty Contest Auctions”. *Marketing Science* 35 (1), pp. 27–54.

Online Appendix

Contents

- A History of Virginia’s SWaM program** **ii**

- B Robustness checks and assumption validity** **iv**
 - B.1 Does vendor participation depend on the bid discount? iv
 - B.2 Does the bid discount vary over time or across categories? vii
 - B.3 Do vendors have different amounts of information about the bid discount? . . . ix
 - B.4 What if vendors have non-rational beliefs regarding the bid discount? x

- C Computational details for approximating the bidding function** **xii**

A History of Virginia’s SWaM program

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% of state spending was going towards minority- and woman-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, woman-owned, and minority-owned (SWaM) vendors:

1. On July 30, 2004, the Governor’s chief of staff, Bill Leighty, sent an internal memo to all state agencies that established a 40% aspirational goal of state expenditures going to SWaM vendors (Leighty, 2004c). In addition, for auctions under \$5000, buyers would have to solicit at least one bid from a woman- or minority-owned vendor. Furthermore, buyers could chose to break the usual lowest-price-wins rule for these small-award auctions if they awarded the contract to a woman- or minority-owned vendor. Note that this 40% goal is never binding or required; buyers are not forced to meet it, but they are strongly encouraged to do so.
2. On September 27, 2004, Bill Leighty sent a Leadership Communiqué to the heads of all state agencies that allowed them to set aside up to 30% of all discretionary spending for small businesses (Leighty, 2004b). This allows buyers to create set-aside auctions using eVA that are only open to certified small businesses.
3. On November 17, 2004, Bill Leighty sent a Leadership Communiqué to the heads of all state agencies that declared that for the purposes of state procurement, “all certified women-owned and minority-owned businesses are also small business enterprises” (Leighty, 2004a). This change in definitions allowed the state to include women- and minority-owned vendors in the preferred group along with small businesses.
4. On December 13, 2005, Governor Warner issued an executive order that reiterated the fact that woman- and minority-owned vendors count as small businesses for the purposes of state procurement. Furthermore, the executive order allowed buyers to award contracts, regardless of the size of the auction, “to a qualified, reasonably priced, certified small business even if it is other than the lowest bidder” (Warner, 2005). In other words, buyers could now deviate from the usual lowest-price-wins rule if they awarded the contract to a SWaM vendor.

5. On August 10, 2006, in response to legal challenges from both within and outside the government, Governor Tim Kaine issued an executive order that established a gender- and race-neutral small business program by removing woman- and minority-owned vendors from the preferred bidding category (Kaine, 2006). Therefore, only true small businesses could receive preferential treatment upon bidding and were eligible for set-aside auctions; woman- and minority-owned vendors no longer qualified.

B Robustness checks and assumption validity

Like other papers in the structural auctions literature, our model and estimation procedure require a number of simplifying assumptions for tractability. Two of the key assumptions are that vendors' decisions whether to enter each auction does not depend on the bid discount level δ , and that all vendors have the same beliefs regarding the bid discount level δ for all auctions within a given year. In this section, we now explain why these two assumptions are plausible in our data setting, provide some ancillary evidence for their face validity, and discuss how our findings would be affected if the assumptions were violated.

B.1 Does vendor participation depend on the bid discount?

No, vendor participation seems to be exogenous.

There are two competing theories that link the bid discount level δ to a vendor's participation decision:

1. Buyers choose the distribution of the bid discounting level on an annual basis. Vendors observe this distribution, and then choose to enter a particular auction if the expected profits exceed the hassle costs of participating. Under this theory, preferred vendors should enter more often in years with a larger bid discount, since their probability of winning and expected profits rise under such a setting.
2. Vendors decide whether or not to enter specific auctions based on their own availability and capability, not based on the level of bid discounting. Buyers observe the likely entry patterns, and then set the level of bid discounting so that they can meet their goals of awarding a portion of contracts to preferred vendors. Under this theory, if preferred vendors are more likely to enter in a given year, the bid discount will be set lower, since buyers do not have to stretch as far in order to award their contracts to them.

These two theories can each be justified based on the level of entry costs. There is no financial cost to submitting a bid, but if the hassle cost of entering an auction is high, then we would expect vendors to be more judicious about entry and to respond to the changes in the level of bid discounting. On the other hand, if the cost of entry is near zero, then vendors that are capable of fulfilling a given contract should bid on it regardless of the bid discounting level.

In the literature, assumptions regarding entry costs in auctions vary based on the marketplace context. For instance, in some situations bidders must invest money in order to observe a private signal regarding their cost or valuation. This model has been commonly

used in the context of timber auctions and highway auctions (Athey, Coey, and Levin, 2013; Athey, Levin, and Seira, 2011; Krasnokutskaya and Seim, 2011). In those settings, firms must incur significant discovery costs as they examine the physical plot of land corresponding to the timber or highway contracts being auctioned off.

With online auctions, the standard assumption is that entry costs are zero, which yields exogenous participation (Athey and Haile, 2007). For instance, Yoganarasimhan (2016) assumes exogenous participation on behalf of sellers because “there are no acquisition costs (a natural assumption in Internet settings, since sellers do not need to spend significant resources to understand their own valuation).” Along these lines, other papers that assume exogenous participation in online auction settings include Coey et al. (2020), Haruvy and Katok (2013), and Yoganarasimhan (2013).

The choice between these two assumptions could affect the model results substantially, although it is not obvious in which direction the results would change. If entry costs are high and vendor entry depends on the bid discount δ , then increasing the bid discount would both increase the number of preferred vendors and decrease the number of non-preferred vendors who participate. The former would strengthen the effect of the affirmative action program, while the latter would weaken it. The net impact of these two opposing forces would depend on the size of the entry costs relative to the overall profits from participating, whether one group had higher entry costs than the other, and the asymmetry in cost distributions between preferred vs. non-preferred vendors.

These two competing theories described at the beginning of this section yield opposite predictions about the relationship between the bid discount level δ and the entry patterns for preferred vendors: under the first theory, these two parameters should be positively associated, while under the second, they should be negatively associated. Given these diverging predictions, this is now potentially a testable assumption. However, we cannot use the same dataset that we use for the rest of this paper – the reason for this is because our current dataset only covers two years, during which time there was a policy change that affected the composition of bidders. Therefore, any large swings in entry probabilities are likely to be due to the policy change, and not due to smaller fluctuations in δ across years. Therefore, we use an ancillary data set which consists of similar Quick Quote auctions in fiscal years 2008 - 2013. There was no similar policy change during this time window, so we can more cleanly decipher the connection between the variation in entry probabilities and the yearly variation in δ .

Our first step is to estimate the distribution of the bid discount δ_t for years 2008 - 2013. We assume that the discount level is distributed log-normal with scale σ and location μ . Estimates from this model are in table A1. For ease of comparison across years, we fix the

scale parameter of the distribution at $\sigma = 14.28$ so that it equals the coefficient from the 2006/2007 log-normal estimate.

Table A1: Log-Normal Estimates for Bid Discount Level (2008-2013)

	Estimate		Std. Error
2008	-3.57	***	0.68
2009	-5.35	***	0.56
2010	-6.44	***	0.61
2011	-6.73	***	0.68
2012	-17.59	***	1.48
2013	-21.09	***	1.48

NOTE: Estimates are of the location parameter μ . Significance levels: *** (0.1 percent), ** (1 percent), and * (5 percent).

We can now estimate the median δ for each year using the formula $\text{Median}(\delta_t) = \exp(\mu_t)$. Combining these yearly median δ values with the observed entry patterns allows us to estimate a logit model of entry. We include the median δ value and the natural log of auction quantity as explanatory variables.

To account for heterogeneity in participation patterns across contracts, we also include fixed effects for each auction title. In this case, an auction title consists of two things: the NIGP commodity code associated with the contract, plus the title of the contract that was submitted by the buyer. For instance, a common auction title in the data is “38548 Strawberries, Sliced, 6/6.5 LB. USDA Grade A” – the five-digit code 38548 refers to the frozen fruit NIGP code, and the rest of the title specifies exactly what the buyer has specified in their contract title. Including these fixed effects helps us account for the fact that some auction titles generally have more entrants than others.

We define “num. preferred participants” as the number of preferred vendors who choose to bid on a given auction, and the potential number of preferred participants as the number of preferred vendors who ever bid on an auction with the same NIGP commodity code, plus five. Subtracting the former from the latter yields the number of preferred non-participants for any given auction.

Each observation is an auction t with auction title a in year y . This yields the following logit model:

$$\ln \left[\frac{(\text{num. preferred participants})_t}{(\text{num. preferred non-participants})_t} \right] = \beta_{0a} + \beta_1 \ln(\text{quantity}_t) + \beta_2(\text{median } \delta)_y + \varepsilon_t$$

Estimates from this model are in table A2. We find strong evidence in support of the second theory: in years where preferred vendors tend to participate more, the level of small

business bid discounting is lower. If vendors were entering based on the level of δ , we would have seen a positive coefficient for $\text{Median}(\delta_t)$. As a result, our counterfactual simulations are run separately for 2006 and 2007 to account for the cost differences and entrant differences between those two years, but we do not allow for the possibility that vendors might enter or leave the market in response to a change in the bid discount value.

Table A2: Logit Estimates for Preferred Vendor Participation Decisions

	Estimate	Std. Error
Median δ_t	-7.504 ***	1.487
Ln(quantity)	0.018	0.016
(Intercept)	-1.455 **	0.461

NOTE: Estimation also includes 2501 fixed effects to account for differences across auction titles. Significance levels: *** (0.1 percent), ** (1 percent), and * (5 percent).

This analysis informs the modeling decisions in this paper. In our structural estimation, we assume that vendors' entry decisions are not affected by the buyer's choice of bid discount level. Prior research has shown that the assumption regarding vendors' entry behavior has a large effect on the size and direction of the subsequent results (Krasnokutskaya and Seim, 2011). The fact that our dataset has variation in the bid discount level over time allows us to examine which assumption is better supported by the data – this provides a point of difference relative to other papers in the literature that cannot test their entry assumptions (Athey, Coey, and Levin, 2013; Athey, Levin, and Seira, 2011; Krasnokutskaya and Seim, 2011; Marion, 2007).

Although our model assumes that vendor participation in a given auction does not depend on the bid discount δ , we do allow for vendor registration (i.e., the pool of potential bidders) to vary based on the rules regarding which kinds of vendors qualify for preferential treatment. Specifically, we assume that the set of available vendors who can bid on contracts is different in 2006 vs. 2007. This assumption is justified by the differences in vendors we observe across those two years (see figure 1 and figure 6), and it helps account for vendors that may strategically enter or exit the auction platform in response to hearing that the government is changing the rules regarding which vendors qualify for the affirmative action program.

B.2 Does the bid discount vary over time or across categories?

No, the bid discount seems relatively stable across time and across product categories.

Vendors in our model have the same beliefs regarding the bid discount δ for all auctions they bid on during the year. This assumption could be flawed if buyers changed their bid

discount in a predictable way. For instance, one possibility is that buyers keep the bid discount low throughout the year but substantially raise it in the last month in order to meet the governor’s expectation of spending 40% of their money with preferred vendors. If vendors became aware of this, they could include this information when formulating their bid.

We examine this issue by estimating a new version of the bid discount distribution that is time-varying. Recall that the bid discount distribution is modeled as log-normal with location μ and scale σ (see section 5.1). We now update this so that there are μ dummy variables by month. The results of this estimation are in table A3.

Table A3: Log-Normal Estimates for Monthly Bid Discount Level

	Estimate	Std. Error	p-value	
Location μ	-17.053	6.407	0.00778	**
Location μ dummy for Feb	-1.5591	3.8101	0.68239	
Location μ dummy for Mar	3.4736	1.6624	0.03667	*
Location μ dummy for Apr	-2.9088	4.0869	0.47663	
Location μ dummy for May	-0.1536	1.4723	0.91692	
Location μ dummy for Jun	8.2061	3.8018	0.03089	*
Location μ dummy for Jul	-2.6211	1.5094	0.08247	
Location μ dummy for Aug	-2.4163	1.606	0.13244	
Location μ dummy for Sep	4.8867	3.7485	0.19235	
Location μ dummy for Oct	3.3514	3.596	0.35135	
Location μ dummy for Nov	8.712	5.843	0.13596	
Location μ dummy for Dec	4.4594	1.8927	0.01847	*
Location μ dummy for 2007	8.3318	3.8599	0.03089	*
Log-Scale $\ln(\sigma)$	2.6209	0.4564	9.31E-09	***

The dummy variables for March, June, and December are statistically significant, but the others are not. However, even when the estimates are statistically significant, they have negligible practical effect on the outcome of interest. The reason these differences have negligible practical effect is because the bid discount delta always enters the first order conditions as part of the term $(1 + \delta)$. See equation 1 and equation 2 in the paper for additional details. Given these values, that key term $(1 + \delta)$ is approximately the same across every single month; in fact, they are identical up to the 4th decimal point. See table A4 for the monthly values of the median δ , as well as the term $1 + \delta$.

The fact that the bid discount does not vary in consequential ways across months supports our modeling assumption that vendors’ beliefs regarding the bid discount δ are stable across time, within each year.

We also repeat this exercise to see whether vendors’ beliefs about the bid discount level

Table A4: Monthly bid discount estimates

	Median bid discount value (δ)	$1 + \delta$
January	3.93×10^{-8}	1.0000
February	8.26×10^{-9}	1.0000
March	1.27×10^{-6}	1.0000
April	2.14×10^{-9}	1.0000
May	3.37×10^{-8}	1.0000
June	1.44×10^{-4}	1.0001
July	2.86×10^{-9}	1.0000
August	3.50×10^{-9}	1.0000
September	5.20×10^{-6}	1.0000
October	1.12×10^{-6}	1.0000
November	2.39×10^{-4}	1.0002
December	3.40×10^{-6}	1.0000

delta vary across different types of contracts (i.e., different industries). Each contract has an NIGP commodity code associated with it, so we include the top 10 most frequently used NIGP codes as dummy variables to see if there are consistent differences in the bid distribution across different industries. In this case the dummy variables are all insignificant, thereby supporting the assumption that vendors’ beliefs regarding the bid discount level delta are stable across different industries.

B.3 Do vendors have different amounts of information about the bid discount?

No, it is unlikely that specific vendors would have more information about the bid discount compared to others. Vendors have very limited information in this context apart from the laws and public statements that the government announced. This lack of information means that it is reasonable to assume that vendors are similarly (un)informed.

The QuickQuote auction system does not give vendors detailed post-auction information on all the other bids that were submitted for that auction. They do see the winning bid, but they do not see how many other bidders there were, what their prices were, or the preferred vs. non-preferred status of any of the bidders. This means that vendors cannot learn more information about the bid discount simply by participating in more auctions.

Therefore, the only way for a vendor to have additional information regarding the realized bid discount value would be to file a Freedom of Information Act request for a number of auctions, and then to estimate a model similar to what we have done in section 5.1 of the paper. This does not happen in practice – in fact, when we submitted a Freedom of

Information Act request to get the data used in this paper, the staff members processing the request said that nobody have ever previously asked for this data.

B.4 What if vendors have non-rational beliefs regarding the bid discount?

One threat to our modeling strategy would be if vendors believed that the bid discount δ was much higher than it actually is. The key issue is that differences in prices between preferred and non-preferred vendors could be attributed either to a high level of δ or to differences in costs between the two groups, and it is not possible to separate these issues empirically.

The model in this paper assumes that all vendors use the yearly-median bid discount δ when formulating their bids. Since these median values are near-zero, if preferred vendors have much higher prices than non-preferred vendors, this implies that preferred vendors as a group generally have higher costs.

However, an alternative possibility is vendors believe that the bid discount is large. For instance, imagine a hypothetical situation where vendors believed that the bid discount was 0.25. In that hypothetical, this higher bid discount could replace higher costs as the explanation for the observed higher bids of preferred vendors. If preferred vendors and non-preferred vendors had the exact same cost distributions and everyone believed that the bid discount level was high, then that would cause their bid distributions to deviate from each other in a pattern similar to what we document in the 2006 data.

One reason why vendors could believe that the bid discount is large would be if they expected the buyers to pick a bid discount δ that would yield exactly 40% of their contract dollars going towards preferred vendors (i.e., what is the minimum bid discount that would allow the buyer to meet the state’s affirmative action goals). Using our data, we find that would imply a bid discount of 0.09 in the 2006 policy environment and a bid discount of 0 in the 2007 policy environment. However, there is a major caveat here: we only have data from QuickQuote auctions, which are not representative of the full eVA auction platform as a whole. The 40% allocation goal is for each buyer’s full set of procurement purchases, not just their QuickQuote purchases. As a result, assuming that buyers set the bid discount delta to hit 40% specifically within their subset of QuickQuote purchases is likely to be misleading.

Unfortunately, the assumption for “what do vendors think the bid discount is” is not a testable one. Nonetheless, there are a few reasons why the assumption in this paper is both reasonable and preferable over the alternative described above:

1. In general, there is substantial evidence in a number of similar procurement contexts showing that there is a substantial gap in the estimated cost distributions between

different observable groups of vendors (small vs. large, preferred vs. non-preferred; etc.). For instance, Athey, Coey, and Levin (2013) and Athey, Levin, and Seira (2011) find this in the context of California timber contracts, Krasnokutskaya and Seim (2011) and Marion (2007) find this in the context of highway auctions, and Flambard and Perrigne (2006) find this in the context of snow removal contracts.

2. In the specific context of Virginia government purchasing, the affirmative action program was instituted specifically to address the fact that a very small percentage of state contracts were going to small, women-owned, and minority-owned vendors. This lack of competitiveness (pre- affirmative action policy) suggests that their costs were much higher as a group.
3. In auctions for California state purchasing and the US Department of Defense, both buyers and vendors are told what the bid discount is. In the Virginia context, buyers and vendors are told that buyers have the option of buying from a preferred vendor who is not the cheapest option, but they were never told that the government was going to implement any particular discount level. As a result, it is not clear why vendors would settle on a much higher discount value, or which specific value they would settle on. By contrast, the assumption used in this paper is consistent with a boundedly rational expectation on the part of the buyer.
4. Although the hypothetical situation described above would explain the gap in bid distributions that is observed in 2006, it would not explain the much smaller gap that is observed in 2007. If vendors non-rationally believed there was a high bid discount, there should still be a substantial gap in the bid distributions of preferred vs. non-preferred vendors in 2007. However, the results indicate this is not the case (see figure 5).
5. Related to the point above: given that the bid distributions are very close together in 2007, if vendors believed that the bid discount was high, this would imply that preferred vendors actually have lower costs than non-preferred vendors (i.e., that small vendors are more efficient than large vendors). This would be an implausible finding.

C Computational details for approximating the bidding function

We assume that the approximate inverse bid function takes the following form:

$$\begin{aligned}
 y_k(b; \alpha_k, \delta) &= b - \text{markup}(b; \alpha_k, \delta) \\
 \ln(\text{markup}(b; \alpha_k, \delta)) &= \alpha_{k,1} + \alpha_{k,2}b + \alpha_{k,3}b^2 + \alpha_{k,4}b^3 \\
 &\quad + \alpha_{k,5}\delta + \alpha_{k,6}\delta^2 + \alpha_{k,7}\delta^3 \\
 &\quad + \alpha_{k,8}b\delta + \alpha_{k,9}b^2\delta^2 + \alpha_{k,10}b^3\delta^3 \\
 &= \alpha_{k,1} + \sum_{i=2}^4 \alpha_{k,i}b^{i-1} + \sum_{i=5}^7 \alpha_{k,i}\delta^{i-4} + \sum_{i=8}^{10} \alpha_{k,i}(b\delta)^{i-7}
 \end{aligned}$$

where k is the vendor's type (preferred or non-preferred), b is the vendor's bid, and δ is the bid discount level. This model provides an intuitive basis for the vendor's bidding rule: the vendor's bid is equivalent to its cost plus a positive markup, and the markup level depends both on the vendor's project cost and the buyer's level of bid discounting.

Denote the first order conditions (equations 1 and 2) as $Q_k(b; \alpha_k) = 0$. We then create a series of grid points in two dimensions: bid values b between the minimum bid \underline{b} and the maximum bid \bar{b} , and bid discount values δ between 0 and 0.40. Denote $grid_{jl}$ as the grid value representing the j -th bid value and the l -th discount value in the grid. This allows us to approximate the inverse bid function by minimizing a least-squares objective function:

$$\min H(\alpha; b, \delta) = \sum_k n_k \sum_j \sum_l Q_k(grid_{jl}; \alpha_k)^2$$

This yields a least-squares estimator for the polynomial approximation of the inverse bid function. Minimizing the function $H(\cdot)$ yields coefficients for the approximate inverse bid function y_k , which maps bids to costs. To ensure that the approximation is well-behaved and that the boundary conditions hold, we minimize the objective function H subject to the following inequality constraints, for all values of b and δ :

$$\begin{aligned}
 \frac{\partial \text{markup}}{\partial b} &< 0 \text{ for all vendors} \\
 \frac{\partial \text{markup}}{\partial \delta} &> 0 \text{ for preferred vendors} \\
 \frac{\partial \text{markup}}{\partial \delta} &< 0 \text{ for non-preferred vendors}
 \end{aligned}$$

Given the parametric nature of our approximation method, one potential concern is

that the results may not be robust to alternative functional form assumptions. Our method follows from Bajari (2001), who recommends using third-order to fifth-order polynomials. As a robustness check, we re-estimate the counterfactual simulations using third- and fifth-order polynomials, with and without interactions between b and δ . The four comparison models are:

$$\text{(Model 1): } \ln(\text{markup}(b; \alpha_k, \delta)) = \alpha_{k,1} + \sum_{i=2}^4 \alpha_{k,i} b^{i-1} + \sum_{i=5}^7 \alpha_{k,i} \delta^{i-4}$$

$$\text{(Model 2): } \ln(\text{markup}(b; \alpha_k, \delta)) = \alpha_{k,1} + \sum_{i=2}^4 \alpha_{k,i} b^{i-1} + \sum_{i=5}^7 \alpha_{k,i} \delta^{i-4} + \sum_{i=8}^{10} \alpha_{k,i} (b\delta)^{i-7}$$

$$\text{(Model 3): } \ln(\text{markup}(b; \alpha_k, \delta)) = \alpha_{k,1} + \sum_{i=2}^6 \alpha_{k,i} b^{i-1} + \sum_{i=7}^{11} \alpha_{k,i} \delta^{i-6}$$

$$\text{(Model 4): } \ln(\text{markup}(b; \alpha_k, \delta)) = \alpha_{k,1} + \sum_{i=2}^6 \alpha_{k,i} b^{i-1} + \sum_{i=7}^{11} \alpha_{k,i} \delta^{i-6} + \sum_{i=12}^{16} \alpha_{k,i} (b\delta)^{i-11}$$

Table A5 provides a comparison of these four models. In all four cases, the inequality constraints are satisfied for every grid point and inequality evaluation. The substantive results are similar across all four models: under the 2006 rules, the affirmative action program yields cost savings for the buyer (on the order of 8 to 12 percent). Under the 2007 rules, the benefits are noticeably smaller.

Table A5: Comparison of alternative polynomial approximations of the bidding function

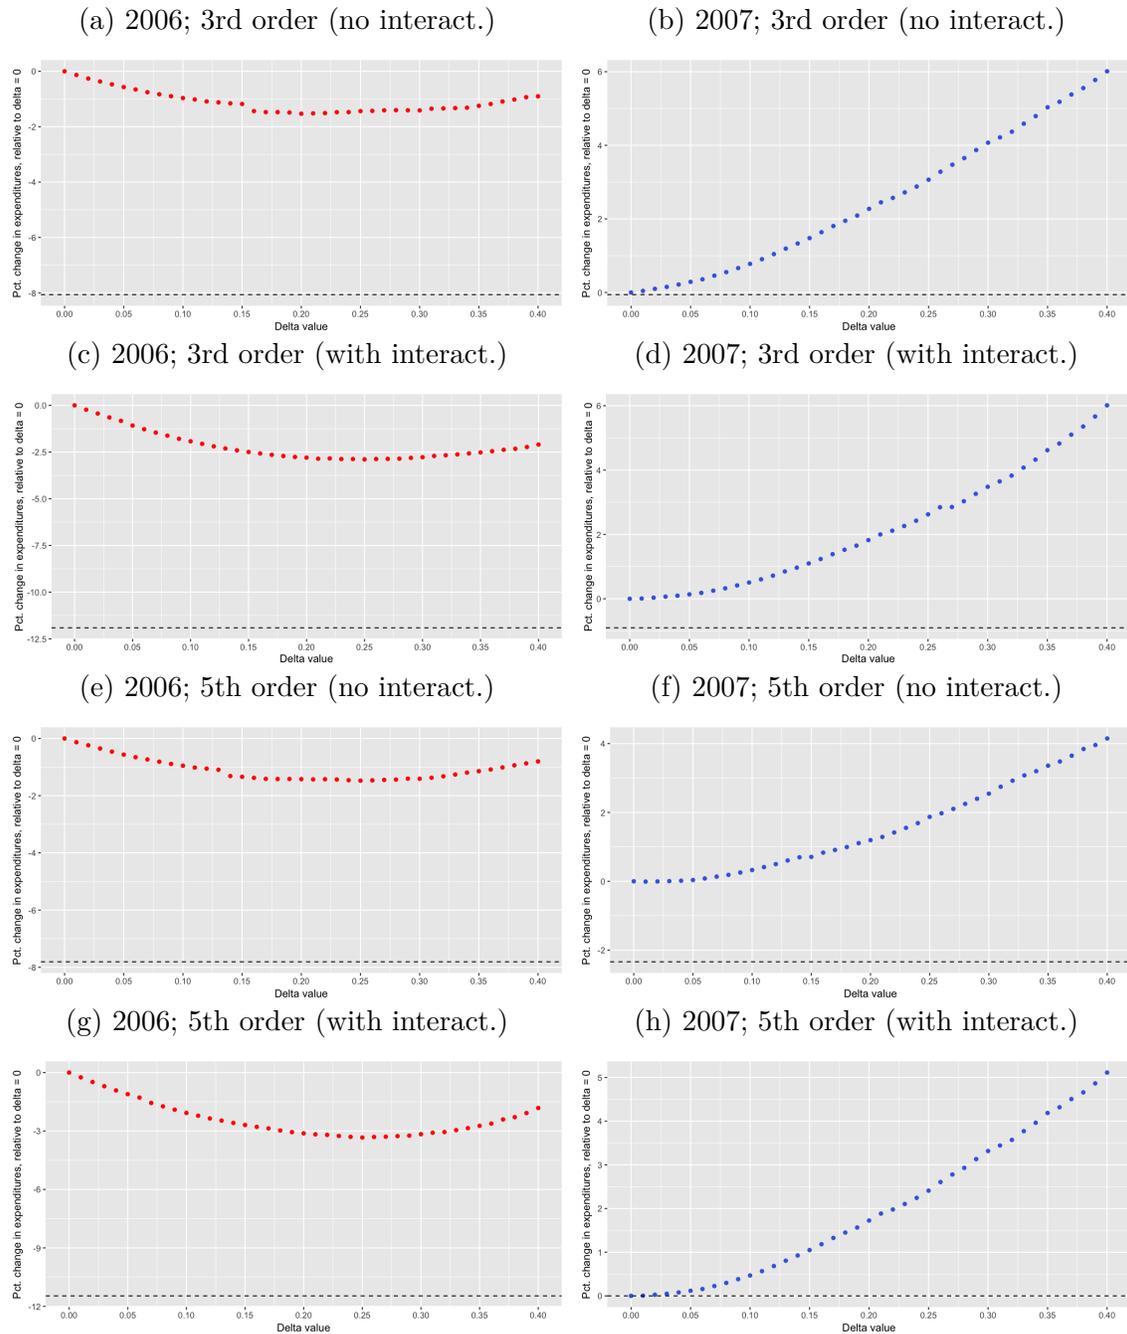
	Model 1	Model 2	Model 3	Model 4
Polynomial order	3rd	3rd	5th	5th
Interactions between b and δ	No	Yes	No	Yes
Num. estimated coefficients (per year)	14	20	22	32
Num. grid points	1640	1640	1640	1640
Num. inequality evaluations (per year)	6560	6560	6560	6560
Num. inequality evaluations that bind	0	0	0	0
Num. inequality evaluations that are violated	0	0	0	0
Pct. cost savings from affirmative action (2006)	8.06	11.91	7.81	11.46
Pct. cost savings from affirmative action (2007)	0.06	0.90	2.34	0.00

Figure A1 compares the different models in terms of their key substantive output: the estimated change in the buyer's expenditures under alternate bid discount values. In the main text, this corresponds to figure 10 for our preferred specification, which is duplicated here as Model 2 (3rd order, with interactions).

Each row of figure A1 corresponds to a different polynomial approximation model from

table A5, and the columns represent the two different years in the data. The key patterns are similar across the four different models. In 2006, there is always a U-shaped pattern between expenditures and the δ level, thereby implying that the buyer would save money by committing to a higher bid discount level relative to the current status quo. However, these benefits disappear in 2007.

Figure A1: Change in expenditures under alternate δ values, by year and model



NOTE: The four rows in the figure correspond to models 1-4 respectively. Expenditure values are calculated relative to the $\delta = 0$ condition. The dashed lines correspond to the first-best outcome in which the buyer optimally varies δ by auction to minimize expenditures.

References

- Athey, S., D. Coey, and J. Levin (2013). “Set-Asides and Subsidies in Auctions”. *American Economic Journal: Microeconomics* 5 (1), pp. 1–27.
- Athey, S. and P. A. Haile (2007). “Nonparametric approaches to auctions”. *Handbook of econometrics* 6, pp. 3847–3965.
- Athey, S., J. Levin, and E. Seira (2011). “Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions”. *Quarterly Journal of Economics* 126 (1), pp. 207–257.
- Bajari, P. (2001). “Comparing Competition and Collusion: A Numerical Approach”. *Economic Theory* 18, pp. 187–205.
- Coey, D., B. J. Larsen, K. Sweeney, and C. Waisman (2020). “Scalable Optimal Online Auctions”. *Stanford University Working Paper*.
- Flambard, V. and I. Perrigne (2006). “Asymmetry in Procurement Auctions: Evidence From Snow Removal Contracts”. *The Economic Journal* 116 (514), pp. 1014–1036.
- Haruvy, E. and E. Katok (2013). “Increasing Revenue by Decreasing Information in Procurement Auctions”. *Production and Operations Management* 22 (1), pp. 19–35.
- Kaine, T. M. (2006). *Commonwealth of Virginia Executive Order 33: Enhancing Opportunities for Small, Women and Minority Owned Businesses*. <http://digitool1.lva.lib.va.us:1801/webclient/DeliveryManager?pid=42799>.
- Krasnokutskaya, E. and K. Seim (2011). “Bid Preference Programs and Participation in Highway Procurement Auctions”. *The American Economic Review* 101 (6), pp. 2653–2686.
- Leighty, W. H. (2004a). *Leadership Communiqué: November 17, 2004*. <http://communique.agencies.virginia.gov/communique47.htm>.
- (2004b). *Leadership Communiqué: September 27, 2004*. <http://communique.agencies.virginia.gov/communique39.htm>.
- (2004c). *Memorandum: SWAM Procurement*. <http://communique.agencies.virginia.gov/memos/swammemo.pdf>.
- Marion, J. (2007). “Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions”. *Journal of Public Economics* 91 (7), pp. 1591–1624.
- MGT of America Inc. (2004). *A Procurement Disparity Study of the Commonwealth of Virginia*.
- 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>.

- Warner, M. R. (2005). *Commonwealth of Virginia Executive Order 103: Promoting Diversity and Equal Opportunity for Small, Woman- and Minority-Owned Business Enterprises in State Procurement*. <http://digitool1.lva.lib.va.us:1801/webclient/DeliveryManager?pid=42168>.
- Yoganarasimhan, H. (2013). “The Value of Reputation in an Online Freelance Marketplace”. *Marketing Science* 32 (6), pp. 860–891.
- (2016). “Estimation of Beauty Contest Auctions”. *Marketing Science* 35 (1), pp. 27–54.