

Bid Exclusion in Public Procurement*

[MOST RECENT VERSION]

Eric Andersen[†]

November 11, 2020

Abstract

Awarding contracts to firms other than the lowest bidder runs directly counter to the intuition behind using auctions in public procurement to keep costs low. In Chicago such bureaucratic discretion — known as bid exclusion — is exercised frequently (in about 14% of auctions) and leads to contracts being awarded to a firm which bid 10% higher in the median case. This mechanical effect increases the cost of public goods and services by 1%. But is this the full cost? This paper (a) documents how Chicago uses bid exclusion as a barrier to entry to receiving public contracts and (b) estimates a simple structural model to capture the full cost of bid exclusion due firms' strategic responses to bid exclusion as practiced by the City. New firms are excluded at a much higher rate than experienced firms (30% v. 8%), as are firms located outside Chicago (17% v. 8%). Since bid exclusion is predictable prior to bidding, sophisticated firms should adjust their bidding strategies in response to the risk that they or their opponents are excluded. Such strategic responses indirectly increase costs across all auctions rather than only in auctions where exclusion occurs. Counterfactual estimates on a subset of auctions suggest the full cost of bid exclusion is at least twice as large as the mechanical increase in cost from awarding contracts to higher bidders. Contracts awarded via bid exclusion are less likely to have cost overruns and have smaller overruns, but the hidden cost of bid exclusion is substantial.

*I am indebted to my advisors Edward Glaeser, Nathan Hendren, Lawrence Katz, and Amanda Pallais. I am also grateful to Andrei Shleifer for his guidance and many helpful conversations. I thank Joshua Hurwitz, Robin Lee, Anthony Yu, and Allen Zhang for helpful comments. Yu Ri Lee provided excellent assistance with the bid tabulation data.

[†]Harvard University, ericandersen@g.harvard.edu

Auctions are a near universal feature of public procurement for the simple reason that they allow governments to award contracts to the lowest cost provider. Given that public procurement accounts for 12% of world GDP, features of auctions that affect costs have potentially enormous efficiency implications (Bosio et al. 2020). One oft-studied class of institutional features that impacts costs is bureaucratic discretion in how the auctions are implemented in practice. Procuring entities possess both well-defined legal powers (e.g., setting reservation prices or awarding contracts via public or non-public bidding) and less well-defined abilities (e.g., manipulating the size of contract or influencing what information is available to firms) that allows them to impact how contracts are awarded and to whom. The form of bureaucratic discretion — bid exclusion — studied in this paper is the ability (and oftentimes the legal mandate) for procuring entities to *disregard* the order of bids for a given contract and award the contract to a higher bidding firm rather than the lowest bidding firm.

Despite being a legal power, when bid exclusion can be exercised is vague at best, making identifying its cost impact difficult. In Chicago — the setting for this paper — procurement officers are tasked with awarding contracts to the “lowest responsive and responsible” bidder.¹ After the bids have been observed, the procurement officer must evaluate whether or not the “winning” firm meets this criteria. If the lowest bidding firm fails to do so, then the next lowest bidding firm is considered and so on. What is not well-defined is what constitutes “responsive and responsible” in practice nor is there public information available for *why* the lowest bidding firm was excluded in any particular auction.

This paper makes the case that bid exclusion impacts the cost of procurement in two ways. The first is the obvious one. By choosing to exclude lower bidding firms in favor of higher bidding firms, the procuring entity incurs a direct cost because the contract award amount is now greater than it otherwise would have been. This cost is salient the procurement officer at the time of the decision because the bids have already been revealed. Furthermore, the direct cost is zero in auctions where the procurement officer awards the contract to the lowest bidding firm. In practice, about 14% of auctions in Chicago during the study period have contracts awarded via exclusion. The total direct cost of this bid exclusion is about \$55 million dollars

¹Municipal Code of Chicago, Chapter 2-92

over all contracts. This represents a large share (10%) of the total amount awarded in these contracts, but a relatively small share (0.8%) of the total amount awarded across all contracts.

The potential second cost of bid exclusion is the unobserved, *indirect* cost due to changes in firms' behavior. Firms which are aware of the possibility their bids or their opponents bids will be set aside by the procuring entity will change their bidding behavior in response. Intuition says firms should increase their bids conditional on costs when bid exclusion is more likely because the probability they will be underbid *and* their lower bidding opponent will be awarded the contract decreases. Therefore, even in auctions where no bids are excluded, there is a potential indirect cost of bid exclusion due to the impact on firms' decisions. This cost cannot be calculated without knowing how firms would have bid when no bid exclusion is possible.

This paper studies the impact of bid exclusion in two parts. The first describes how procurement officers in the City of Chicago implement bid exclusion using data on City contracts and bid tabulations. In the 14% of auctions where exclusion occurs, the median contract is awarded to a firm that bid 10% higher than the lowest bid in the auction. This is a salient *direct cost* of bid exclusion when the City excludes bids. While the legal motivation for bid exclusion is to evaluate bids and firms on a case-by-case basis for whether they can complete the contract to the specifications expected by the City, the decisions made by the procurement officers reflect a distinct pattern. The City uses bid exclusion to erect an additional barrier to entry for certain types of firms. New firms are excluded at a much higher rate than experienced firms (30% v. 8%), and firms from outside Chicago are excluded at a much higher rate than firms from within Chicago (17% v. 8%). In order for firms to adjust their strategies in anticipation of bid exclusion, exclusion needs to be predictable prior to firms submitting their bids for a particular auction.

The second part of the paper quantifies the unobserved, indirect cost of bid exclusion caused by firms adjusting their bidding strategies in response to the risk that they or their opponents will be excluded. This paper builds a simple structural model of bid exclusion which builds on the canonical symmetric independent private values (IPV) reverse auction model. In this model firms take the probability of exclusion as exogenous. Firms choose their optimal bid by trading off their profit conditional on being awarded

the contract and their probability of being awarded the contract, which in the case of bid exclusion is different from the probability of being the lowest bidder. Using bid tabulations from the auctions run by Chicago’s DPS and the firms identified in the city’s repository of contracts, this paper estimates the distribution of latent cost as specified by the model. Armed with this cost distribution and estimates of each bid’s cost, the standard symmetric IPV bidding strategy yields what bid each firm would have submitted in the absence of the probability of exclusion ($p = 0$).

Using this method to compare each observed bids to its counterfactual, the indirect impacts of bid exclusion are large relative to the naïve calculation. I find that the indirect costs of exclusion even in a small ($n = 350$) subset of auctions is 20% larger than the total direct cost across auctions in the full sample.

The remainder of the paper is organized as follows. [Section 1](#) gives an overview of the key institutional features of public procurement in the City of Chicago and describes the data. [Section 2](#) describes the practice of bid exclusion in Chicago. [Section 3](#) introduces a model of public procurement auctions with bid exclusion. [Section 4](#) estimates the model. [Section 5](#) leverages the estimated model to construct firms’ counterfactual bids and calculates the indirect cost of bid exclusion. [Section 6](#) concludes.

1 Empirical Setting and Data

1.1 Public Procurement in Chicago

The Department of Procurement Services (DPS) administers all public contracts in the City of Chicago. The lifecycle of a typical contract proceeds as follows. First, a User Department within the city government, e.g. the Chicago Department of Transportation or the Department of Police, identifies a good or service that requires outside contracting and contacts DPS. Next, procurement officers within DPS work with the User Department to finalize an in-house specification that describes the goods and/or services required and decides what method of procurement will be used to fulfill the need.

Not every contract awarded by DPS is awarded via auction. By state and municipal law, contracts above a fixed threshold must be subject to a pub-

lic and competitive process. This threshold was \$10,000 prior to December 2004 and increased to its current level of \$100,000. The form of this competitive process falls under the discretion of DPS but takes two primary forms: first-price sealed-bid auctions and Requests for Proposals (RFPs). Auctions are used when the DPS has a clear understanding of the exact goods and/or services desired and “cost is the sole priority.”² DPS uses RFPs when it is willing consider more than only the cost when deciding which firm is contracted. A typical example of a contract awarded via auction would be for repaving a road, for being the official supplier of office supplies to the City for two years, or for the purchase of new police cruisers. A typical example of a contract awarded via RFP (or the closely related RFQ) would be for wiring City buildings with broadband Internet where the City relies on the expertise and insight of the applicants to help inform “how much” work it wishes to undertake even if that means increasing the cost substantially. For contracts below the legal threshold, DPS can choose to award the contract via a non-public (but ideally still competitive) process such as soliciting “bids” privately via telephone.

If the DPS decides to bring a contract to auction, a written specification is publicly announced and a date and time is set as the deadline for accepting bids from vendors. DPS maintains a list of registered vendors, which is used to verify firm characteristics — such as being minority- or women-owned — that impact various bid incentives. Any firm, however, can submit a bid packet for any auction. Depending on the complexity of the project — major construction, for example — DPS holds meetings to thoroughly describe the details of the specification to interested firms. Firms who send representatives to these meetings are also contacted with updates to the specification. Firms submit sealed bid packets to the Bid and Bond Room prior to the posted auction date and time. At the announced time, the bid packets are unsealed and read aloud. Once the winning bidder is announced, the winner is usually bonded using a fixed proportion of their bid.

A major assumption of the model in [Section 3](#) is that the number of firms is exogenous and public information. This is a strong assumption and many procurement models incorporate endogenous entry of firms. Why not in this context? While it is true that not all construction firms participate in all con-

²This this comes directly from DPS, but obviously the existence of bid exclusion would imply cost is never the sole priority but perhaps the primary one.

struction auctions, for example, DPS requires firms that take specification packets to identify their firm (so that they may receive any and all up-to-date changes to the specification) and then posts this information publicly online. This, combined with the fact that firms attend joint informational meetings about major projects, suggests that firms likely have a good sense of how many and which firms specifically are likely to enter the auction.

1.1.1 Exclusion

DPS has a legal obligation to award contracts to the “lowest, responsible” bidder in a given auction. The use of this term gives DPS the bureaucratic discretion to skip over bids that are low but deemed not “responsible.” What exactly constitutes “responsible” is not defined statutorily in either the Illinois State Code or the Chicago Municipal Code. The Illinois State Code (65 ILCS 5/8-10) defines a responsible bidders as “[a] person who has the capability in all respects to perform full the contract requirements and the integrity and reliability that will assure good faith performance.” The Chicago Municipal code is equally as vague: “[I]f the lowest responsive and responsible bidder is unable or unwilling to perform the work, or if such bidder’s performance is insufficient for the city’s needs, the work, or any uncompleted portion, shall be performed by the next lowest responsive and responsible bidder”

1.2 Data Sources

1.2.1 Contract Awards.

Contract data going back to the early 1990s — both initial contracts and all subsequent modifications — are publicly available on the [City of Chicago Data Website](#) which was launched in 2011 at the beginning of Mayor Rahm Emanuel’s administration. This data is updated on a daily basis and extends back to contracts dated in 1993. Coverage in the early years is incomplete, but grows over time. By the early 2000s, there is near universal coverage which corresponds to an overhaul of the computer database system used by DPS to track and update contracts. Each contract in the data identifies the firm by name and alphanumeric ID number and a variety of characteristics

about the commodity or work under contract.³ The two crucial pieces of information are the specification number — which links the contract to the DPS document describing the goods and services under contract — and the awarded value of the contract. In addition to the contract data, the City of Chicago also releases payment information for each city contract going back to 2002. Payments diverge from the awarded value of the contracts due to contract positive or negative modifications or if the City requires more or less than the amount predicted under the signed term agreement.

1.2.2 Bid Tabulations.

DPS maintains a [searchable database](#) of the universe of information on the auctions and submitted bids used to award contracts going back to early 2004. Auctions prior to March 2004 are available in PDFs going back to early 2000. Each bid tabulation provides the specification number of the DPS document detailing the goods and services up for auction. This is the key linkage used to merge the Contract data and the Bid Tabulations data. Additional information on the end date and time of the auction, User Department, etc are also available. The bid tabulations provide the full list of firms that submitted bid packets in the auction and their bid amounts. In rare cases, bids are rejected outright (<1% of all bids). Unfortunately, when bids are rejected the submitted bid amount is not available and there is usually no explanation provided for why the bid packet was deemed unacceptable. They could have been bids that were too high, too low, or perhaps the bid submission demonstrated the firm could not meet the expectations set in the auction's specification regardless of the price. The focus of this paper is on bids that are excluded, not rejected, and so these bids are dropped.

1.2.3 Variable Definitions

Firm Location. Each contract in the data contains a contact mailing address for the firm which allows for firms to be categorized by whether or not they are located in Chicago. For firms with multiple addresses — implying they have offices at multiple locations or moved sometime during the sample period — I identify them as being located in Chicago if they ever held a contract address in Chicago.

³The firm's address is also always available, which will become relevant later.

Firm Experience. Firm experience is measured on the date of each auction. I use two closely related measures of firm experience. The primary one is the number of contracts held by the firm prior to the auction date which is constructed using the award dates in the auctions data. The secondary measure is the number of prior auctions the firm has participated in (regardless of whether or not the firm wins the auction).

Excluded Bids. Excluded bids are not explicitly identified in the bid tabulations but is inferred from the contracts data. [Table 1](#) provides an example bid tabulation where bid exclusion occurs. The key linkage between the contracts data which identifies the winner of the auction and the bid tabulations is the specification number. In this case, Firm 2 is the firm identified in the contracts data as holding the contract that corresponds to the specification number used in the auction. Since Firm 1 has a lower bid than Firm 2 yet did not receive the contract, the bid is identified as excluded.

Payment Overruns. I supplement the contracts data with payment data also available through the City of Chicago. This data is available post-2002. I define a payment under/overrun as the difference between the awarded amount and the ultimate expenditure made by the City during the life of the contract.

1.2.4 Analysis Sample

The raw sample of bid tabulations contains about 34,000 accepted, non-zero bids.⁴ I make three restrictions on the contracts and bid tabulations to construct the analysis sample. The first step is to exclude certain auction formats that make it difficult to match to the contract information. The contract award amounts are for the total expected amount of the project, but some auctions have bids at the price-level. Other auctions have bids at an “adjustment factor” level where firms bid at ratios — 1.23 for example — of a reference price book that DPS provides. The next step is to restrict to auctions where the specification can be matched to a contract and the winning

⁴Bids can be zero. This occurs when there are multiple auctions for a single specification in which a firm bids in one auction but not another. For example, the purchasing of rock salt by CDOT falls under a single specification document but the auctions are done at the CDOT facility-level. Firms submit bids that are zero on facilities where they do not wish to win a contract.

bid can be identified. Not all auctions can be successfully matched to contracts because some auctions are aggregated under a single contract. The third restriction is to restrict to auctions where the award amount for the contract matches the bid submitted by the firm who is awarded the contract. The motivation for this restriction is to avoid auctions where the contract amount is heavily negotiated after the bid submission process. This is the subset of auctions where it is clear the most salient factor in determining the award of the contract. After all the restrictions, 1,254 auctions remain in the analysis sample.

1.2.5 Summary Statistics

Table 2 provides counts and summary statistics for the the bids in the analysis sample. The upper panel shows statistics calculated over all bids in the sample, and the lower panel restricts the bids to only the lowest bids in each auction. Roughly 15% of the lowest bids are excluded. The proportion of lowest bids that are from Chicago or are from new firms are slightly lower than the full sample of bids.

2 The Practice of Bid Exclusion in Chicago

This section documents several empirical facts about bid exclusion and explores to what extent bid exclusion is a predictable feature of the auctions.

2.1 Direct Costs of Exclusion

In a given auction, DPS observes the bids as submitted by the firms and must then decide whether or not to exclude the lowest bidding firm in favor of a higher bidding firm. I define the direct cost of bid exclusion in auction t as

$$(\text{Direct Cost})_t = (\text{Bid Awarded Contract})_t - (\text{Lowest Bid})_t$$

If DPS awards the contract to the lowest bidding firm, then $(\text{Direct Cost})_t = 0$. If DPS decides to exclude the lowest bid, then the direct cost is strictly positive, $(\text{Direct Cost})_t > 0$.⁵ This cost is direct in the sense that it is the im-

⁵There is of course the knife-edge case where two firms bid the exact same amount that is lower than all other bids submitted in this auction. While possible, since I do not observe

mediate consequence of bid exclusion and is the one most salient to both the procurement officers making the decision and the economist observing the data. Whatever the exact objective function of the DPS procurement officers — differences in expected work quality, differences in time delays, direct bribery or extortion, etc. — they should exclude the lowest bid when the total change in monetized benefits from switching firms outweighs this direct (change in) cost.

The direct costs of bid exclusion across all auctions can be easily calculated from the data. [Table 3](#) provides summary statistics of the direct costs for contracts awarded via exclusion. The first column shows summary statistics across all auctions. It suggests the direct costs of bid exclusion are large and varied. The total direct cost is \$55.8 million with an average direct cost of roughly \$311,000. The dispersion is large, however, with some extremely large observations and a median direct cost of \$48,263.75. The total amount awarded in these 180 contracts is about \$560 million, meaning bid exclusion accounts for 10% of the awarded cost of these contracts.⁶ The total amount awarded across all contracts in the analysis sample is \$6.1 billion meaning the direct cost of excluding bids adds a little less than 1% to the total amount awarded.

Generally speaking, excluded bids can be divided into two groups: clear mistakes and close bids. Clear mistakes occur when the lowest bidding firm deviates wildly from the rest. In these cases, it is not clear the direct cost is really a true cost from the perspective of the City. If firms mistakenly submit bids lower than they should have — due to inattention or misunderstanding the City’s contract specification — then it is unlikely the City would truly pay what that firm bid. The firm would potentially have to renegotiate for a high payment ex post. How would this bias the distribution of direct costs though? Consider a simple example of an auction with four bids: \$10,000, \$110,000, \$115,000, and \$120,000. The City excludes the lowest bid. The calculated direct cost from above is \$100,000 but it is unlikely that the lowest bidding firm could have completed the contract for \$10,000 without requiring extra payments. The fact that the other three firms provides valuable information about what the true cost is likely to be. The true direct cost would be

this event in my sample, it makes sense to think of exclusion as imposing a strict increase in costs to the City.

⁶This ignores ex-post modifications to the contract award amounts which will affect the total amount paid, not awarded.

lower in this case. The second group of excluded bids are ones that are much closer where the lowest bidding firm could have completed the contract for the amount they bid but the City excludes them anyway.

The second column of [Table 3](#) excludes auctions where the lowest bidding firm bids less than 50% of the second lowest bidding firm, a rough proxy for the lowest bidding firm making a mistake.⁷ This hopefully eliminates the majority of the mistaken low bids and can give more informative estimates of the moments of the direct cost distribution. The median direct cost is about 33% lower as is the total direct cost across auctions, but the pattern is strikingly similar to that in the full sample.

2.2 Firm Characteristics

In order for firms to anticipate the likelihood of being excluded or having their opponents excluded, the rate of exclusion must be predictable based on observable characteristics of the firms. This section addresses whether this holds for two salient features of every firm: where the firm is located and how experienced the firm is in Chicago’s public procurement system.

2.2.1 Firm Location Predicts Exclusion

[Figure 1](#) Panel A plots the relationship between firm location and the rate of exclusion. Local firms have a massive advantage over outside firms in the likelihood of being awarded a contract conditional on being the lowest bidder. Firms from outside Chicago are excluded 17.5% of the time while firms from within Chicago are only excluded 8.5% of the time. Auction characteristics may explain the difference in exclusion rates, however, if local firms participate in auctions that are less likely to be awarded via exclusion for reasons unrelated to firm characteristics. To address this issue [Figure 1](#) Panel B shows the probability of exclusion after residualizing the exclusion outcome on several auction-level characteristics, namely indicators for the contract’s

⁷This difference is known as the bid spread.

$$(\text{Bid Spread})_t \equiv (\text{Second Lowest Bid})_t - (\text{Lowest Bid})_t$$

Since the DPS procurement officers do not pre-commit to excluding bids the size of the bid spread may have an effect on whether or not they choose to exclude the lowest bid. For example, it is more likely the benefits from exclusion outweigh the direct cost when $(\text{Bid Spread})_t$ is small.

User Department, indicators for the type of contract (e.g., commodities, construction, consulting, etc.), and indicators for the number of participating firms. This narrows the difference in exclusion rates by about 50% to 6 percentage points.

The fact that Chicago-based firms have a lower probability of being excluded has several potential explanations. The first is that the City of Chicago has a preference for local firms that is unrelated to the quality of the product/work. Local preferences are explicit in other aspects of the city's procurement policies. For example, the Chicago Municipal Code establishes local hiring quotas on construction contracts (so that even outside firms that win contracts must have a certain fraction of worker-hours worked by Chicago residents). It would be natural for this preference to be reflected in how exclusion is applied in practice. The City might simply just prefer to "buy local." A less benign explanation, of course, would point to Chicago's long history of political corruption and a myriad of criminal convictions involving local contracts. Local firms may have closer ties to local officials who exert influence over contracts, and the City is less willing to exclude them all else equal.

2.2.2 Experience Predicts Exclusion

Firms' connections to Chicago procurement have a chronological dimension in addition to their location. Firms that have participated in auctions or have prior contracts are known to DPS whereas new firms are not. DPS may exhibit the same sort of preferences for such known firms.

Figure 2 plots how the lowest bidder's experience affects the rate of exclusion. The left panel plots a bincatter of the rate of exclusion as a function of the number of contracts the lowest bidding firm has been awarded by the City prior to the date in which the auction occurs. The lowest bidding firms are first divided into two groups: new firms and firms with prior experience. Firms with prior experience are then divided into quartiles based on the number of contracts they have received. A striking discontinuous "L" shape is readily apparent. Almost one third of lowest bidders with no experience are excluded which jumps to only 12% of firms in the low quartile of the experienced firms which have only received 1 or 2 contracts. The rate of exclusion declines slightly as firms become more experienced but these

differences are not statistically significant.

The right panel of [Figure 2](#) plots the same relationship but with a slightly different definition of the lowest firm's experience. In this case, experience is measured as the number of auctions the lowest bidding firm participated in prior to the date of that auction regardless of whether or not the firm received contracts in those auctions. The same negative relationship between the rate of exclusion and firm experience emerges except now the change is considerably more continuous. The difference between the rates of exclusion for the lower quartiles of the (positive) experience distribution is explained by the fact that firms never one a contract are excluded multiple times before exiting or receiving a contract.

2.2.3 Experience or Location?

The sections above demonstrate that DPS has predictable differences in their decisions to exclude firms along two dimensions examined individually. The fact that firms from outside Chicago tend to be less experienced than firms located from Chicago, however, suggests one difference may be biased due to the omission of the other. [Figure 3](#) plots the rates of exclusion across the four groups defined by the interaction between the lowest firm's location and experienced (zero prior contracts and at least one prior contracts). Three observations are apparent. First, the rates of exclusion are ordered how one would expect. The type of firm most at risk for exclusion is a new firm from outside Chicago (26%), and the type of firm least at risk is an experienced firm from Chicago (8%). Second, within each experience level Chicago firms are excluded less often but in neither case is this difference significant. Third, within each location, experienced firms are excluded less often than new firms with no prior contracts. This difference is not significant for Chicago firms but is large and significant for firms located outside Chicago.

[Figure 4](#) approaches the same question but allows for more variation in firm experience by showing the same relationship as in [Figure 2](#) but now conditioning on firm location. Here there seems to be little difference between experienced firms from within and outside Chicago, but inexperienced firms are excluded at a slightly higher rate than new firms, the same observation as before.

2.2.4 Firms and Contract Outcomes

If these firm characteristics predict rates of exclusion, it is natural to wonder whether the same characteristics also predict contract outcomes that DPS *should* (hopefully) be concerned with. This is a tough question to answer because even if firm entry were exogenous, the practice of bid exclusion means the firms that “survive” are heavily selected across some (potentially unobserved) dimension. If firm characteristics are no predictive of contract outcomes, that is consistent with (but not definitive evidence of) the hypothesis that DPS is effectively screening out bad firms.

The contract outcome observable in the data is payments. Public contracts are incomplete and a long literature the impact of renegotiation due to changes in the contract requirements, unexpected problems, etc. Payment outcomes are defined relative to the contract’s award mount. I split these into three outcomes. The first is when the payment is less than the award amount. In other words, the firm completed the contract for less than it bid. The second is when the payment is equal to the award amount.⁸ The final payment outcome is when the payment amount is greater than the award amount by more than 0.1% of the award amount. This is a “cost overrun” which I interpret as a negative contract outcome.

Contract modifications in Chicago procurement are not unilateral decisions on the part of the firms. Instead they are the product of ex-post negotiations between the firm and the procurement officer who oversees the contract. This similar to other types of contract renegotiation common within procurement departments.

Figure 5 shows the relationship between the *winning* firm’s location and the payment outcome of the contract. Contracts are divided into two groups based on the location of the firm who received the contract and then the share of contracts that have each payment are shown. This is in fact the 2-by-3 frequency table of payment outcomes by firm location. The awarded firm’s location does not appear to be predictive of the payment outcome. This is confirmed by using a Chi-square test of independence which fails to reject the null that the payment outcomes are independent of firm location ($p = 0.25$).

⁸I allow for some leeway because payments can be rounded slightly to the nearest \$100 or so. As long as the payment is within 0.1% of the award amount on either side, I categorize them as equal.

Figure 6 plots the relationship between the *winning* firm's experience level and the contract's payment outcome. Whereas firm's location did not appear to predict payment outcomes, what is immediately clear is the same general relationship between the rate of exclusion and experience as shown in Figure 2 is present in the relationship between the contract outcome and experience. Less experience firms who received contract awards had larger payment overruns than more experienced firms.

2.3 Exclusion and Contract Outcomes

The last issue to consider is whether contracts awarded via exclusion have better payment outcomes than contracts awarded to the lowest bidder. Again if DPS is using bid exclusion to exclude "irresponsible" firms (and their bids), then contracts awarded via exclusion should have outcomes at least as good as those awarded in the standard manner.

Luckily this appears to be the case. Figure 7 plots the relationship between a contract's payment outcome and its award mechanism – either to the lowest bidder or via exclusion. Under either mechanism, the contract has a payment less than the award amount about 40% of the time. Contracts awarded via exclusion, however, are twice as likely to have a payment exactly equal to the amount awarded which is the amount the winning firm bid. This difference comes entirely from a reduction in the likelihood of having a payment overrun, which is a little over 10 percentage points less likely to occur for contracts awarded via exclusion.

Figure 8 plots the distribution of payment overrun magnitudes conditional on having a payment overrun. A stark pattern emerges. The distribution of overruns for contracts awarded to the lowest bidder has a large concentration of "small" overruns – nearly 50% – that are in the first bin. The share of contracts awarded via exclusion in the same bin, however, is less than half that rate. The share of contracts in the "middle" bins is then relatively constant across both groups. In the final bin, for contracts with a payment overrun greater than 50% of the award amount, the pattern flips and contracts awarded via exclusion are now twice as likely to fall in this bucket as contracts awarded to the lowest bidder.

This pattern has several potential explanations worth mentioning. At the low end of the distribution, remember that firms who win contracts via

exclusion have bids – and therefore award amounts – that are larger than the the lowest bidders. In the median case, the “winning” bid is 10% larger than the lowest bid. One possible story is that firms who bid lowest anticipate being able to negotiate some contract modifications with DPS that allow the payments to slight overrun the award amount. DPS, knowing that it is paying less than what it would have to the second lowest bidder, still pays less relative to exclusion and renegotiates for small overruns. If DPS has awarded the contract to a higher bidder, however, it is now paying more than it otherwise would.

2.4 Ex Post Bid Spreads

The decision to exclude a given firm’s bid is made after the bids have been revealed, so while firms make their strategic decisions based on the probability their opponents will bid a certain way, DPS makes its exclusion decision in each auction taking the bids as fixed. Therefore, the probability a given bid amount — say \$100,000 — is excluded could depend on the other bids in the auction — for example, whether the next lowest bid is \$110,000 or \$1 million.

To understand the impact of other bids on the probability the lowest bid is excluded, [Figure 9](#) plots the probability of exclusion against the auction’s bid spread. The bid spread is defined the difference between the lowest bid and the second lowest bid either in dollar terms or as a share of the lowest bid. The relationship between exclusion and the bid spread is U-shaped. This indicates two types of bid exclusion: bids that are excluded because (a) they are much lower than the next lowest bid (who is the likely recipient of the contract if the lowest bid is excluded) or (b) they are just below the next lowest bid. The benefit of excluding the first type of lowest bid is the same as the illustrative example above. It is likely the low bid is a mistake rather than a strategic decision on the firm’s part and does not reflect the final cost of the contract. The benefit of excluding the second type of lowest bid is that DPS can possibly incorporate its beliefs about the relative quality of the firms without increasing cost significantly. If a low quality firm has underbid a high quality firm by 1%, the increase in quality is more likely to outweigh the increase in cost than if the low quality firm underbid by 15%. Hence, close auctions have a higher probability of DPS excluding lowest

bids.

3 A Model of Procurement with Bid Exclusion

This section constructs a model of procurement auctions with the possibility that the lowest bidder will be excluded by the procuring entity. The model builds on a standard sealed-bid, first-price reverse auction that is common in the literature on procurement auctions started by Milgrom and Weber (1982).

3.1 Setup and Notation

Auctions. Consider a set of identical auctions $t = 1, 2, \dots, T_N$ where $N > 1$ firms bid for a government contract with the City. The City announces the specification for auction t publicly, gives firms time to prepare sealed bid packets, and then closes bidding at an announced date and time.

Firms. Firms are risk-neutral and profit-maximizing. Each firm is one of two types which differ only in the probability of being excluded conditional on being the lowest bidder, denoted p_1 and p_2 . The number of firms of each type — denoted N_1 and N_2 — participating in each auction t is exogenous and public information prior to the auction. Firms are indexed $i = 1, 2, \dots, N$ within each auction.

Costs. The cost for firm i to complete the contract in auction t is a random variable C_{it} . This cost is identically and independently distributed according to a well-behaved distribution $F : [\underline{c}, \bar{c}] \rightarrow [0, 1]$ which has a compact support. This distribution is public information to all firms and the City. The realization of each firm's cost is denoted c_{it} and is private information known only to firm i prior to the bids being unsealed.

Payoffs. After observing its private cost c_{it} for the auction, firm i submits a sealed bid b_{it} to the City. After all firms have submitted their bids, the City reveals all N bids and awards the contract to the lowest non-excluded firm.

Firm i 's profits are

$$\pi_{it} = \begin{cases} b_{it} - c_{it} & \text{if firm } i \text{ is awarded contract} \\ 0 & \text{otherwise} \end{cases}$$

Contract Awards and Exclusion. The key feature of this model is the possible exclusion of the lowest bid and the differentiation between submitting the lowest bid and being awarded the contract. After revealing the bids and before awarding the contract, the government flips a weighted coin. With probability p , the lowest bid will be excluded, and the government awards the contract to the second lowest bid. With probability $1 - p$, the lowest bid is accepted and the contract is awarded to the lowest bidder as in a standard first-price sealed-bid auction.

Equilibrium Strategies Notation. The solution concept for this game is a symmetric Bayesian Nash equilibrium. The equilibrium strategy maps each firm's private cost onto the bid submitted to the City.

$$\beta : [\underline{c}, \bar{c}] \rightarrow [\underline{b}, \bar{c}]$$

The standard assumptions that (a) the bidding strategy is strictly increasing in cost, $\beta' > 0$, and (b) the firm that draws the highest cost \bar{c} will bid \bar{c} hold. I will denote the inverse bidding strategy as $\eta = \beta^{-1}$ which is strictly increasing in costs. Therefore, in equilibrium,

$$\beta(c_{it}) = b_{it} \quad \text{and} \quad \eta(b_{it}) = c_{it}$$

3.2 Firm Problem

Being the lowest bidder and being awarded the contract are the same event in a standard first price auction. Bid exclusion decouples these events. If a firm is the lowest bidder, then it still has to survive exclusion to win the contract. If a firm is the second lowest bidder, then there is a non-zero probability the opponent who bid lower than it did is thrown out and the firm wins the contract. In the model, risk-neutral firms choose bids to maximize the expected payoff and tradeoff increasing their profits conditional on receiving the contract with decreasing their probability of receiving it. The

only innovation in the model presented here is to augment the award probability to incorporate this difference in probabilities.

$$\Pr(\text{awarded contract} | b_{it}, \eta_1, \eta_2) \neq \Pr(\text{lowest bidder} | b_{it}, \eta_1, \eta_2)$$

Knowing the Under the risk of bid exclusion, a firm's optimal bidding strategy is the solution to the optimization problem

$$\beta(c_{it}) = \operatorname{argmax}_{b_{it}} \left\{ (b_{it} - c_{it}) \times \Pr(\text{awarded contract} | b_{it}, \eta_1, \eta_2) \right\}$$

When it is convenient, I suppress the exogenous parameters (p_1, p_2, N_1, N_2) for the sake of clarity but all strategies are still conditioned on them.

3.2.1 Probability of Contract Award

Without loss of generality, assume firm i is a type 1 firm. Under the model, there are three events that result in firm i being awarded the contract with a bid b_{it} :

- (i) firm i is the lowest bidding firm *and* firm i is not excluded by the City
- (ii) exactly one type 1 firm underbids firm i *and* that firm is excluded by the City, and
- (iii) exactly one type 2 firm underbids firm i *and* that firm is excluded by the City.

Holding its opponents strategies fixed, if firm i submits bid b_{it} then the probability the firm is awarded the contract is therefore

$$\begin{aligned} \Pr(\text{awarded contract} | b_{it}, \eta_1, \eta_2) = & \\ & (1 - p_1) [1 - F(\eta_1(b_{it}))]^{N_1 - 1} [1 - F(\eta_2(b_{it}))]^{N_2} \\ & + p_1 (N_1 - 1) F(\eta_1(b_{it})) [1 - F(\eta_1(b_{it}))]^{N_1 - 2} [1 - F(\eta_2(b_{it}))]^{N_2} \\ & + p_2 N_2 F(\eta_2(b_{it})) [1 - F(\eta_1(b_{it}))]^{N_1 - 1} [1 - F(\eta_2(b_{it}))]^{N_2 - 1} \end{aligned}$$

3.3 Profit Maximizing Markups

The necessary first-order condition of this optimization problem can be rearranged to express the firm i 's markup as

$$\underbrace{b_{it} - c_{it}}_{\text{markup}} = \left[\alpha_1(b_{it}) \times \underbrace{\frac{(N_1 - 1) f(\eta_1(b_{it})) \eta_1'(b_{it})}{1 - F(\eta_1(b_{it}))}}_{\text{standard IPV term}} + \alpha_2(b_{it}) \times \underbrace{\frac{N_2 f(\eta_2(b_{it})) \eta_2'(b_{it})}{1 - F(\eta_2(b_{it}))}}_{\text{standard IPV term}} \right]^{-1} \quad (1)$$

where

$$\alpha_1(b_{it}) \equiv \frac{1 - 2p_1 + p_2 N_2 \frac{F(\eta_2(b_{it}))}{1 - F(\eta_2(b_{it}))} + \mathbb{1}(N_1 \geq 2) \times p_1 (N_1 - 2) \frac{F(\eta_1(b_{it}))}{1 - F(\eta_1(b_{it}))}}{1 - p_1 + p_2 N_2 \frac{F(\eta_2(b_{it}))}{1 - F(\eta_2(b_{it}))} + p_1 (N_1 - 1) \frac{F(\eta_1(b_{it}))}{1 - F(\eta_1(b_{it}))}}$$

and

$$\alpha_2(b_{it}) \equiv \frac{1 - p_1 - p_2 + p_1 (N_1 - 1) \frac{F(\eta_1(b_{it}))}{1 - F(\eta_1(b_{it}))} + \mathbb{1}(N_2 \geq 1) \times p_2 (N_2 - 1) \frac{F(\eta_2(b_{it}))}{1 - F(\eta_2(b_{it}))}}{1 - p_1 + p_1 (N_1 - 1) \frac{F(\eta_1(b_{it}))}{1 - F(\eta_1(b_{it}))} + p_2 N_2 \frac{F(\eta_2(b_{it}))}{1 - F(\eta_2(b_{it}))}}$$

The left-hand side of this equation is the firm i 's profit markup above cost conditional on being awarded the contract. The right-hand side consists of two types of terms: exclusion terms $\alpha_1(b_{it})$ and $\alpha_2(b_{it})$ and two standard IPV-esque terms. Under loose assumptions on the probabilities of exclusion, these exclusion terms are positive and bounded above by 1.⁹

$$0 \leq p_1 < \frac{1}{2} \implies 0 < \alpha_1(b_{it}) \leq 1$$

$$0 \leq p_1 + p_2 < 1 \implies 0 < \alpha_2(b_{it}) \leq 1$$

Intuitively, the bounds on the exclusion terms suggest the “distorted” markups due to exclusion are greater than the markups under no exclusion (the standard IPV model). To understand the extent to which this holds requires solving for the optimal bidding strategies under both scenarios: exclusion and non-exclusion. In practice, however, this has more to do with model specification. the standard methodology used to estimate these

⁹In practice, these conditions are easily met the estimated probabilities of exclusion for both types are less than $\frac{1}{3}$.

models replaces the standard IPV terms with estimated bid distributions – which are fixed regardless of model used. Hence if the true model is the one described, then estimating the model in the standard way will lead to lower estimated markups due to misspecification.

3.3.1 No Exclusion Case

A special case of this model is the case of no exclusion, where the lowest bidder is guaranteed the contract. This has two implications. First, there is no “distortion” due to exclusion.

$$p_1 = p_2 = 0 \iff \alpha_1(b_{it}) = \alpha_2(b_{it}) = 1$$

Second, firms are now symmetric and thus have identical bidding strategies:

$$\eta(b_{it}) \equiv \eta_1(b_{it}) = \eta_2(b_{it})$$

Both implications immediately imply the markup characterized in Equation (1) collapses to

$$b_{it} - c_{it} = \frac{1 - F(\eta(b_{it}))}{(N_1 + N_2 - 1) \times F(\eta(b_{it})) \times \eta'(b_{it})} \quad (2)$$

Solving from the firm’s optimization problem yields an equivalent condition. This is exactly the first-order condition in the symmetric IPV procurement model with N bidders.

4 Structural Estimation

This section describes the methodology used to estimate the latent costs \hat{c}_{it} that corresponds to every observed bid b_{it} which are then used to create an estimate \hat{F} of the latent cost distribution F faced by firms. First, I address the clear heterogeneity across auctions by homogenizing the bids. Second, I estimate the distribution of these homogenized bids, holding firms’ bidding strategies fixed. Third, I use the necessary condition for firms’ profit maximizing bids (as under the model specification) to invert the homogenized bid onto the implied “pseudo-cost” to the firm. Finally, I estimate the

distribution of these imputed “pseudo-costs” using standard kernel density estimation.

4.1 Creating “Homogenized” Bids

To estimate the model using existing off-the-shelf methods requires data on homogenous auctions. This section follows Krasnokutskaya (2011) and Haile and Kitamura (2018) in using heterogeneity across auctions that enters in a way that is multiplicatively separable from firms’ idiosyncratic costs.¹⁰

4.1.1 Cost Structure

Following Krasnokutskaya (2011) and Haile and Kitamura (2018), I make the following two assumptions. The first is on the cost structure which I assume to have the form

$$C_{it} = \Gamma(X_t) \times C_{it}^h$$

As before C_{it} is the cost to firm i to complete the contract in auction t . This cost is assumed to have two components, an auction-specific component $\Gamma(X_t)$ that is an (unknown) transformation of a vector of observable auction characteristics X_t . The second is the firm i ’s idiosyncratic cost that is drawn i.i.d. from a well-behaved distribution F .¹¹

The second assumption is that the auction characteristics X_t are independent of the joint distribution of firm-specific components

$$X_t \perp\!\!\!\perp (C_{1t}, C_{2t}, \dots, C_{Nt})$$

Under these two assumptions, Krasnokutskaya (2011) shows that if

$$\beta_{it}^h(C_{it}|N, F_c)$$

¹⁰The same literature discusses heterogeneity that enters in a multiplicatively additive way. See Athey and Haile (2007). See Bajari, Houghton, and Tadelis (2014) for a practical implementation to highway contracts.

¹¹This cost component is sometimes referred to as the “homogenized” cost.

is the equilibrium strategy for the index case where $\Gamma(X_t) = 1$ then

$$\beta_{it}(C_{it}|X_{it}, N, F_C) = \Gamma(X_t) \times \beta_{it}^h(C_{it}|N, F_C)$$

is the equilibrium bidding strategy when $\Gamma(X_t) \neq 1$.

4.1.2 Estimating Auction Components

In this application, the auction-specific shock will be estimated by estimating auction fixed effects in a simple OLS framework. I will therefore replace $\Gamma(X_t)$ with Γ_t . It is straight forward to estimate the fixed effects by fitting the OLS specification

$$\ln b_{it} = \gamma_t + u_{it}$$

Having estimated the fixed effects $\hat{\gamma}_t$, I transform the observed bids into the homogenized observed bids.

$$b_{it}^h \equiv \frac{b_{it}}{\exp(\hat{\gamma}_t)}$$

Before moving onto estimating the normalized bidding strategy, one practical issue in how the auction fixed effects γ_t are estimated is worth discussing. That is the sensitivity of the fixed effects estimates to bids that demonstrate clear mistakes on the part of bidding firms and are not reflective of firms strategic behavior. For example, consider an auction where the three submitted bids are \$10,000, \$110,000, \$120,000. The estimated auction-specific shock would be about \$51,000 which indicates that one firm received an extremely low idiosyncratic shock while the other two firms received extremely high idiosyncratic shocks. Whether or not this is an accurate estimate of the shock depends on whether the first firm's bid is a reflection of their strategy or the simple mistake of placing the decimal in the wrong place. This situation is in fact one of the reasons the discretion to exclude the lowest bid is valuable to the procuring entity. If the first firm meant to bid \$100,000, there is little chance they would be able to complete the project for \$10,000 in a satisfactory way without renegotiation, delays, etc.

My solution to this issue is to drop the "mistaken" bids prior to esti-

imating the auction fixed effects. In the above example, the auction-specific shock estimated with only the two high bids is about \$115,000 and this would be the value used to compute the two homogenized bids. The true auction-specific shock would be \$109,000, so this is still a biased estimate but the bias is greatly reduced compared to using the \$51,000 estimate. The obvious candidates for these mistakes are the highest and lowest bids in each auction that are extreme outliers relative to the other bids in the auction. To identify which are the most likely to be mistakes I take the ratio of the highest bid to the second highest bid and the ratio of the lowest bid and the second lowest bid. I then drop the bids with the highest and lowest ratios, respectively. In practice, I drop few bids at this stage. My stopping criteria is once the model no longer predicts wildly negative estimated costs for bids the city found acceptable and awarded the contract.

4.2 Inverting the Homogenized Bids

4.2.1 Type-Specific Probabilities of Exclusion

The steps up until this point have not required specifying which firms belong to which types. In order to preserve the variation across types and avoid splitting the sample into too small type-1-by-type-2 bins, I group firms based on whether or not the firm has ever received a contract from the City of Chicago regardless of their location. Type 1 and type 2 firms are simply New or Experienced firms. Given this division of firms, the exogenous probability that a firm of each type will be excluded conditional on being the lowest bid is

$$\begin{aligned}\hat{p}_{\text{Firms With No Prior Contracts}} &= 0.302 \\ \hat{p}_{\text{Firms With 1+ Prior Contracts}} &= 0.086\end{aligned}$$

See [Figure 10](#) which shows the estimated parameters for each group. These parameters — along with the number of firms of each type — are observable prior to bidding and taken as exogenous by firms.

4.2.2 Estimated Bid Distributions

Now that the homogenized bids b_{it}^h have been calculated, I estimate separate bid distributions, one for each cell of auctions defined by the number of participating bidders of each type. This follows directly from the arguments made by Guerre, Perrigne, and Vuong (2000) and the subsequent literature it ignited. When the number of firms of each type (N_1, N_2) is fixed, then the bidding strategies (η_1, η_2) are also fixed. Since the (inverse) bidding strategies are also strictly monotonic, the distribution of bids is the same as the distribution of costs. For type 1 firms for instance,

$$\begin{aligned} G_1(b_{it}^h) &= F(\eta_1(b_{it}^h)) \\ g_1(b_{it}^h) &= f(\eta_1(b_{it}^h)) \times \eta_1'(b_{it}^h) \end{aligned}$$

The left-hand side of both equations above can be estimated from the observed homogenized bids in a fixed N_1 -by- N_2 cell of auctions using the empirical cumulative density function (ECDF) and a kernel density estimate (KDE).¹² The same exercised can be done for type 2 firms.

4.2.3 Creating Pseudo-Costs

With the estimated distributions of bids of each type \hat{G}_1, \hat{G}_2 and their kernel density estimates \hat{g}_1, \hat{g}_2 , the homogenized costs c_{it}^h for each firm in each auction can be calculated by rearranging Equation (1) and substituting in the estimated bid distributions for the latent cost distributions and bidding strategies.

$$\hat{c}_{it}^h = b_{it}^h + \left[\hat{\alpha}_1(b_{it}^h) \times \frac{(N_1 - 1) \hat{g}_1(b_{it}^h)}{1 - \hat{G}_1(b_{it}^h)} + \hat{\alpha}_2(b_{it}^h) \times \frac{N_2 \hat{g}_2(b_{it}^h)}{1 - \hat{G}_2(b_{it}^h)} \right]^{-1}$$

The right-hand side of this equation is observable and has been estimated. Hence an estimated homogenized pseudo-cost \hat{c}_{it}^h can be calculated for each observed homogenized bid b_{it}^h . The kernel density estimates \hat{g}_1 and \hat{g}_2 , however, are biased near the edges of the observed homogenized bid support for each type. This bias will in turn bias the estimated pseudo-costs \hat{c}_{it}^h for

¹²Following recommendations from Guerre, Perrigne, and Vuong (2000), I use a tri-weight kernel with a bandwidth selected according to Silverman's Rule of Thumb, separately for each type with each cell of auctions.

homogenized bids near b_{it}^h the boundaries. In line with the standard recommendations in the literature, I therefore disregard the pseudo-costs that lie within one bandwidth of the minimum and maximum bids of either type.

4.2.4 Identification

Identification of the pseudo-costs \hat{c}_{it}^h requires the relationship between the homogenized bids and pseudo-costs to be strictly increasing for each of the two firm types within each N_1 -by- N_2 auction cell. [Figure 11](#) shows this relationship for the auction cell with two new firms and two experienced firms and confirms the costs are identified by inverting the bids. This figure plots only the bids that lie within the trimmed support as recommended by Guerre, Perrigne, and Vuong (2000). In general, the costs are identified in auction cells where the total number of firms of any type is strictly greater than two, $N_1 + N_2 > 2$.

This identification fails in the three auction cells where the total numbers of firms of each type is two, $N_1 + N_2 = 2$. [Figure 12](#) plots the relationship between bids and estimated costs for the auctions with two new firms. When $N_1 + N_2 = 2$, one or both of the exclusion terms α_1 and α_2 become unbounded. At the high end of the bid distribution, this distortion can cause the relationship between bids and estimated costs to be non-monotonic. As a result, I cannot use the resulting estimated pseudo-costs to generate counterfactual bids as described in [Section 5](#). I therefore drop all the auctions where only two firms participate.

5 Measuring Indirect Costs

This section calculates the indirect cost of bid exclusion due to the change in bidding strategies of sophisticated firms.

5.1 Counterfactual Bids

This section uses the estimated distribution of latent costs \hat{F} to calculate bids under the counterfactual scenario where the firms do not respond to the probability of exclusion. In practice this is straightforward. As noted above, when $p_1 = p_2 = 0$ the model collapses to the commonly used symmetric

independent private values (IPV) reverse auction model. The differential equation defined in the no exclusion case has a known, closed-form solution.

$$\beta_{\text{IPV}}^h(c_{it}^h | p_1 = p_2 = 0, N_1, N_2, F) = c_{it}^h + \frac{\int_{c_{it}^h}^{\bar{c}} [1 - F(s)]^{N_1 + N_2 - 1} ds}{[1 - F(c_{it}^h)]^{N_1 + N_2 - 1}} \quad (3)$$

For each bid with a valid pseudo-cost c_{it}^h , the counterfactual IPV bid under no exclusion can be estimated by substituting c_{it}^h and \hat{F} into [Equation \(3\)](#).

$$\hat{\beta}_{it}^{h, \text{IPV}} \equiv \beta_{\text{IPV}}^h(\hat{c}_{it}^h | p_1 = p_2 = 0, N_1, N_2, \hat{F})$$

which can be compared do the observed homogenized bid b_{it}^h .

5.2 Impact of Exclusion

Now that the counterfactual bids under the no exclusion $\hat{\beta}_{it}^{h, \text{IPV}}$ regime have been estimated, I can determine the indirect cost due to bid exclusion by comparing the observed homogenized bids b_{it}^h to their respective counterfactuals. I define the change in bids due to bid exclusion as

$$\Delta_{it} = b_{it}^h - \hat{\beta}_{it}^{h, \text{IPV}}$$

The changes due to bid exclusion for a representative auction are shown in [Figure 13](#). This figure shows three bidding strategies. Two are the “observed” bidding strategies for new and experienced firms, relating their observed homogenized bid with their estimated pseudo-costs. Conditional on cost, the experienced firms bid higher than the new firm. The third bidding strategy is the counterfactual bidding strategy $\hat{\beta}_{it}^{h, \text{IPV}}$. Both new and experienced firms have the same optimal bidding strategy conditional on cost in this counterfactual scenario because they are treated the same by the procurement officer and there is no probability of exclusion. The impact of exclusion Δ_{it} — represented by the vertical distance between each firms bids — is clearly positive for both types of firms, but is larger for experienced firms.

5.2.1 Decomposing Total Costs

All firms change their bids in response to the possibility exclusion, but change in costs is only the difference between the observed award amount and the counterfactual award amount in the auction with no exclusion. Let firm i be the firm awarded the contract in auction t in the data with or without exclusion. Let firm j be the firm that would have been awarded the contract in auction t in the counterfactual scenario. That is, firm j has the lowest IPV bid:

$$\hat{\beta}_{jt}^{h, \text{IPV}} < \hat{\beta}_{kt}^{h, \text{IPV}} \quad \forall k \neq j$$

The total change in cost in auction t due to bid exclusion is the difference $b_{it}^h - \hat{\beta}_{jt}^{h, \text{IPV}}$. This total change can be decomposed into the direct cost of exclusion and the indirect cost of exclusion in the following manner:

$$\begin{aligned} \overbrace{b_{it}^h - \hat{\beta}_{jt}^{h, \text{IPV}}}^{\text{(Total Exclusion Cost)}_t} &= \overbrace{(b_{it}^h - b_{jt}^h) \times \mathbb{1}(b_{it}^h \geq b_{jt}^h)}^{\text{(Direct Cost)}_t} \\ &\quad + \underbrace{(b_{it}^h - b_{jt}^h) \times \mathbb{1}(b_{it}^h < b_{jt}^h) + (b_{jt}^h - \hat{\beta}_{jt}^{h, \text{IPV}})}_{\text{(Indirect Cost)}_t} \end{aligned}$$

Each auction falls under one of three possible cases. [Table 4](#) summarizes the cost components. The first case is when the same firm would have submitted the lowest bid in both auctions and the procuring entity chooses not to exclude that firm. In this case, $i = j$ and $(\text{Direct Cost})_t = 0$. The total cost of exclusion is due to the change in the bidding strategy of the winning firm and is therefore unobserved. The second case is when the procuring entity does not exclude the lowest observed bid, but the identity of the lowest bidding firm is changed due to differences in optimal bidding strategies, i.e. $i \neq j$. How could this occur? Referring back to [Figure 13](#), note that Experienced firms have a larger change in bidding strategies Δ_{it} than New firms because it is more likely a New firm will be excluded conditional on submitting the lowest bid. The normal monotonicity of bidding strategies is not necessarily preserved when the two lowest bidding firms are of different types. It is possible for an Experienced firm to have a slightly smaller cost c_{it}^h than a New firm but then overbid that New firm in response to bid

exclusion. In this case, the indirect cost is still positive but is slightly smaller than if the “original” winner firm j had received the contract. The final case is the case of bid exclusion, $b_{it}^h > b_{jt}^h$. Firm j would have submitted a lower bid than Firm i both with and without exclusion, but in the real world the procuring entity intervened, creating a positive cost. This is the only case in which the direct cost is positive.

5.3 Aggregate Indirect Costs

Table 5 shows summary statistics for the subset of auctions with valid indirect cost estimates. It is immediately clear that the trimming of the pseudo-costs as described in Section 4.2.3 has a meaningful impact on the number of auctions for which the indirect cost can be calculated. Only 351 of the auctions in the original analysis sample have valid estimates for the counterfactual IPV bid. The rest are either auctions with only two firms or have a homogenized bid that lies too close to the boundary of the bid support.

Despite the magnitude of the indirect costs of exclusion relative to the direct costs is quite large. The direct costs of exclusion in Table 3 totals \$57.6 millions across all 1,254 auctions in the analysis sample, and the indirect costs for only slightly over 25% of the auctions almost \$70 million. The median auction in the post-estimation sample has a lowest bid that is 2.35% higher than the lowest bid that would be submitted in an IPV auction. This cost is several times that of the total direct cost of bid exclusion which is about 0.8% of all contracts awarded. It is clear that the majority of the costs of bid exclusion are hidden.

6 Conclusion

Auctions are widely used by governments to ensure contracts go to the lowest bidder and deliver goods and services without overpaying. The procuring entity, however, may wish to award the contract to a firm other than the lowest bidding firm after observing the firm’s identities and the amount they bid. This generates an obvious cost — contract awards must increase mechanically — but the full cost of bid exclusion is unknown because firms likely respond strategically to the risk that they or their opponents will be excluded.

This paper first studies bid exclusion in Chicago over 20 years to understand when the City's procurement officers exclude certain bids and firms and whether the City's preferences are predictable. The second part of the paper studies how firms should respond to the City's preferences for exclusion. The paper builds a simple structural model of first prices auctions with an exogenous probability of exclusion. Firms belong to one of two types, with different levels of risk. I estimate this model on New and Experienced firms and generate counterfactual bids to measure to full cost of bid exclusion. Using the counterfactual bids, I find that while the direct costs of exclusion represent less than 1% of the total contract award amounts over the sample period, the indirect costs are several times that and the total cost of bid exclusion is at least 3%.

References

- Athey, Susan, and Philip A Haile. 2007. "Chapter 60 Nonparametric Approaches to Auctions." In *Handbook of Econometrics*, edited by James J Heckman and Edward E Leamer, 6:3847–3965. Elsevier.
- Bajari, Patrick, Stephanie Houghton, and Steven Tadelis. 2014. "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs." *The American Economic Review* 104 (4): 1288–1319.
- Bosio, Erica, Simeon Djankov, Edward L Glaeser, and Andrei Shleifer. 2020. "Public Procurement in Law and Practice." w27188. National Bureau of Economic Research.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica: Journal of the Econometric Society* 68 (3): 525–74.
- Haile, Philip, and Yuichi Kitamura. 2018. "Unobserved Heterogeneity in Auctions."
- Krasnokutskaya, Elena. 2011. "Identification and Estimation of Auction Models with Unobserved Heterogeneity." *The Review of Economic Studies* 78 (1): 293–327.
- Milgrom, Paul R, and Robert J Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica: Journal of the Econometric Society* 50 (5): 1089–1122.

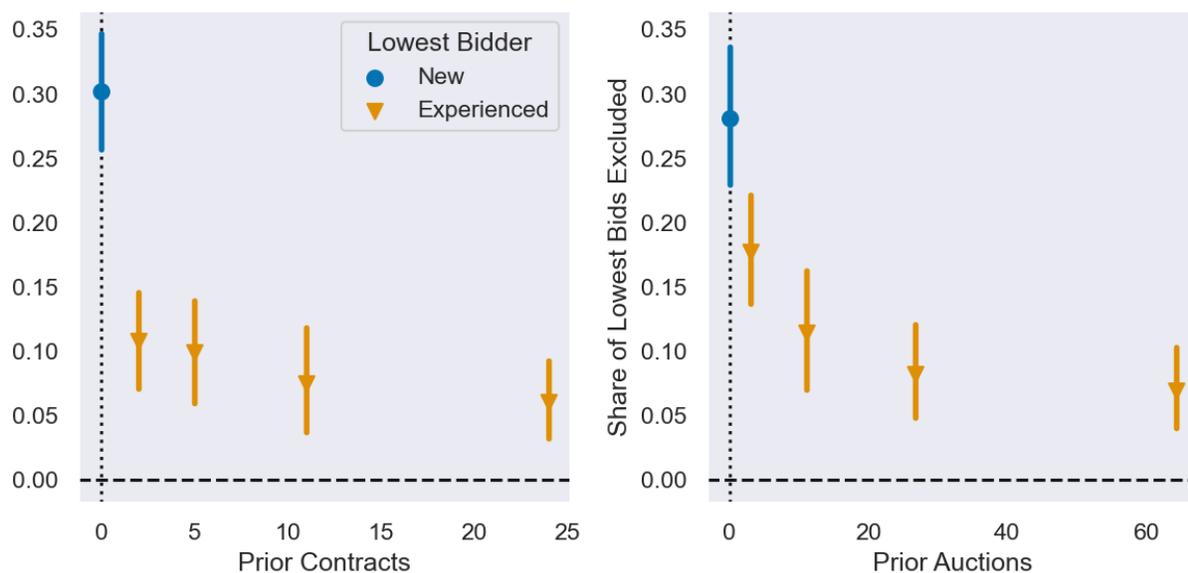
Figures

Figure 1: Firm Location and Exclusion Probability



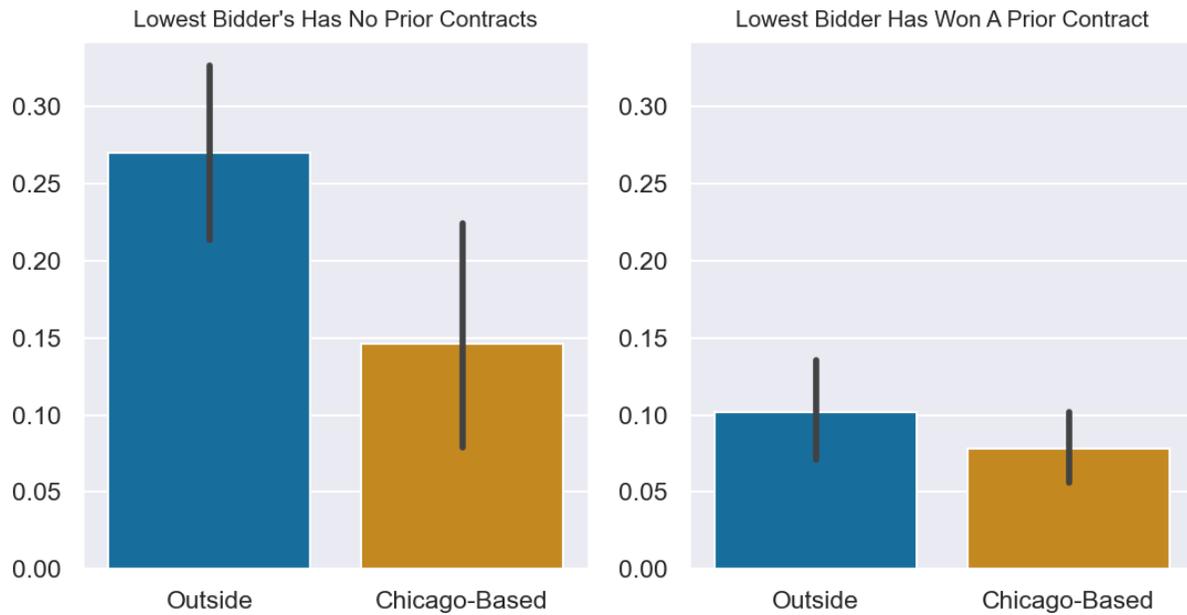
Notes — This figure shows how the rate of bid exclusion depends on the location of the lowest bidding firm. In Panel A, the rate of exclusion for firms located outside Chicago is nearly double that of “home” firms. Panel B residualizes an indicator for exclusion in auction-level characteristics to control for the composition of auctions that firms from each location enter. See [Section 2.2.1](#) for details.

Figure 2: Firm Experience and Exclusion



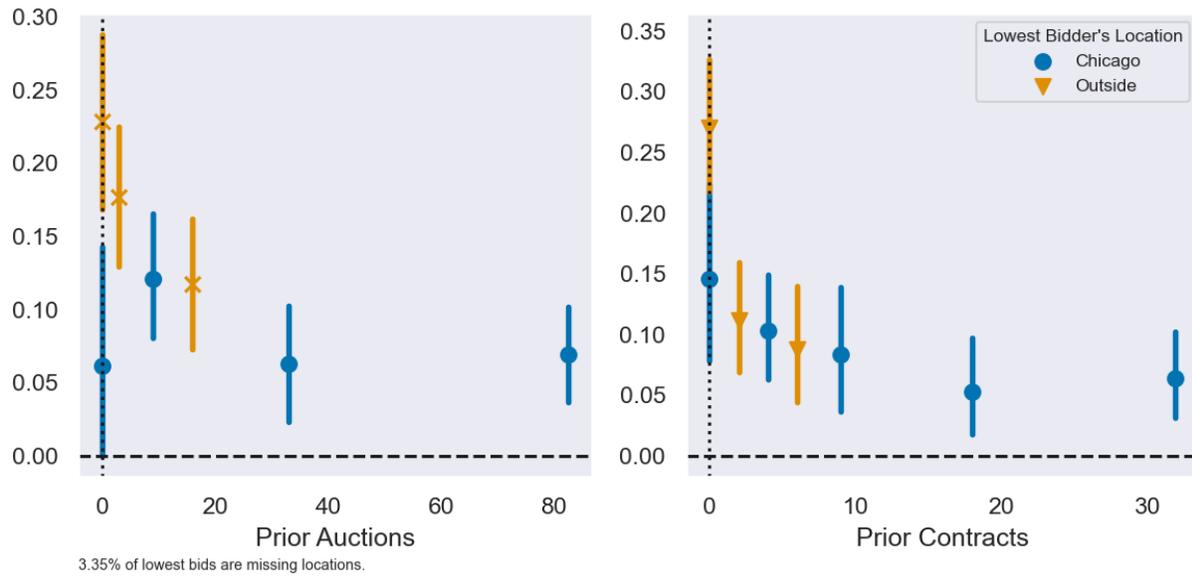
Notes — This figure shows a binscatter of the relationship between the rate of exclusion and the experience level of the lowest bidding firm. The y-axis is the share of auctions where the lowest bidder is excluded. The x-axis is a measure of the lowest bidding firm’s experience. In Panel A this is the number of contracts the firm has received prior to entering the auction. In Panel B this is the total number of auctions the firm has entered prior to entering the auction. Within each panel the firms are broken out into two groups. The first is new firms that have an experience level of zero. The second is experienced firms are those with strictly positive levels of experience. See [Section 2.2.2](#) for details.

Figure 3: Rate of Exclusion, By Experience Level and Location



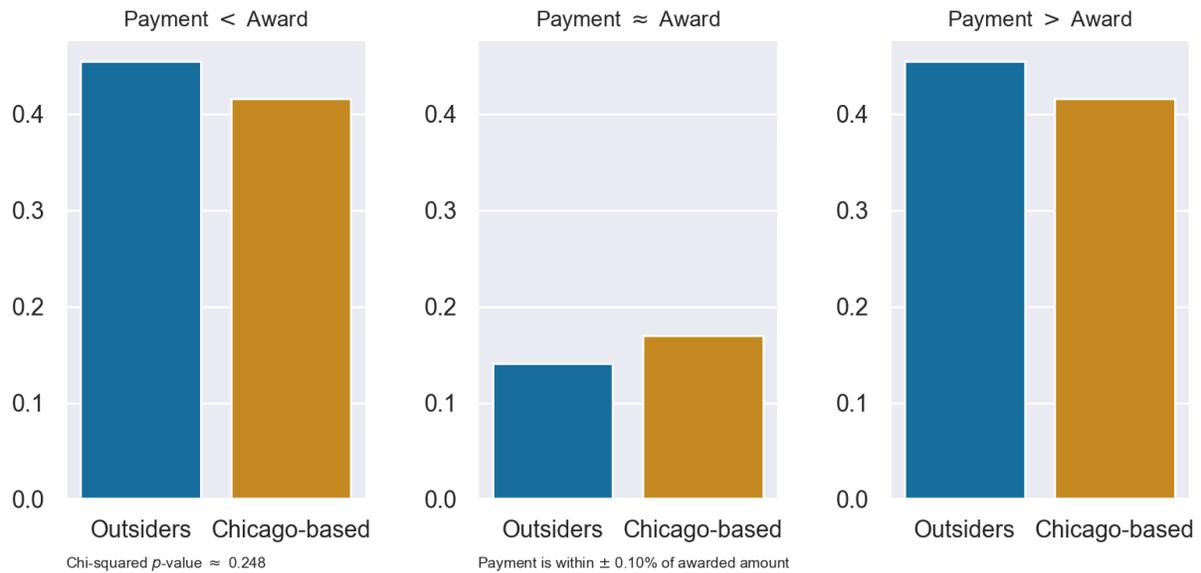
Notes — This figure displays the interaction between the lowest bidding firm's location and experience level in predicting the rate of exclusion. The left panel shows how the rate of exclusion among new firms varies by location. The right panel shows how the rate of exclusion among experienced firms varies by location. See [Section 2.2.3](#) for details.

Figure 4: Rate of Exclusion By Experience and Location



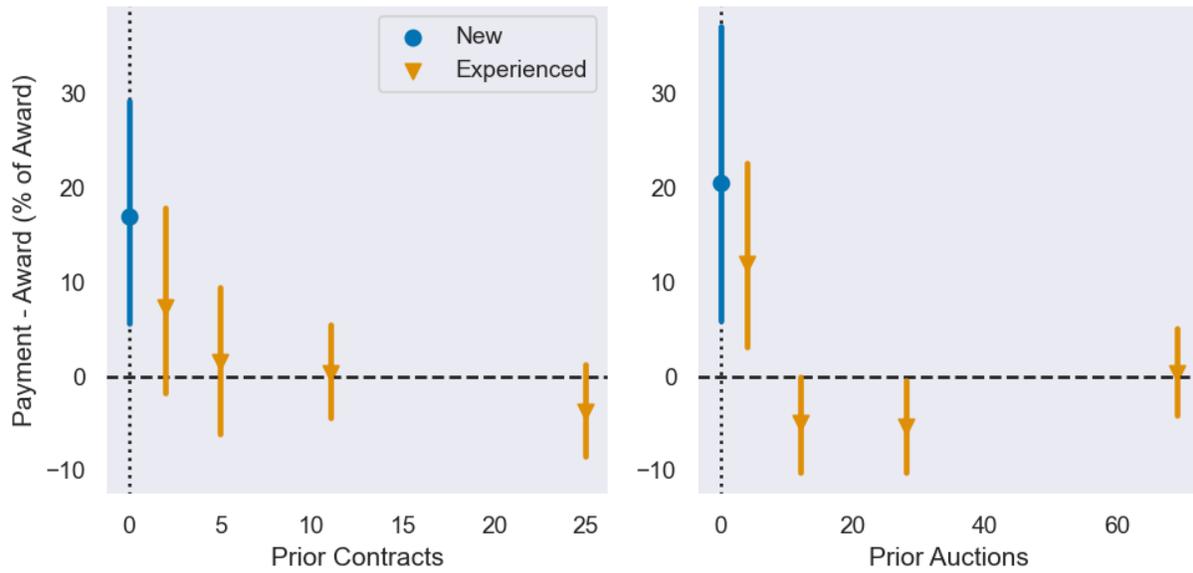
Notes — This figure plots the relationship between the rate of exclusion and the firms experience level, by location. This is the same relationship as shown in [Figure 2](#) but now firms are broken out by location. See [Section 2.2.3](#) for details.

Figure 5: Payment Outcomes and the Winning Firm's Location



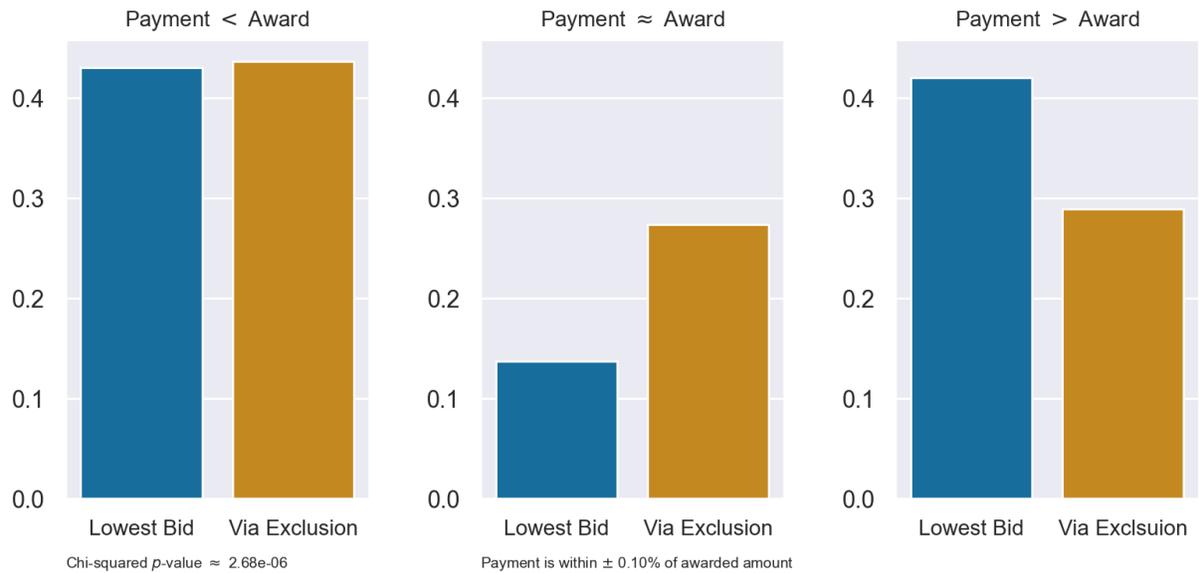
Notes — This figure shows the relationship between the contract's cost outcome and the winning firm's location. Contracts can have three outcomes: payment less than the contract award, payment approximately equal to the contract award, and payment in excess of the contract award (a cost overrun). The reported p -value is from a Chi-squared test of independence for a 2-by-3 frequency table of contract counts. See [Section 2.2.4](#) for details.

Figure 6: Payment Outcomes and the Winning Firm's Experience



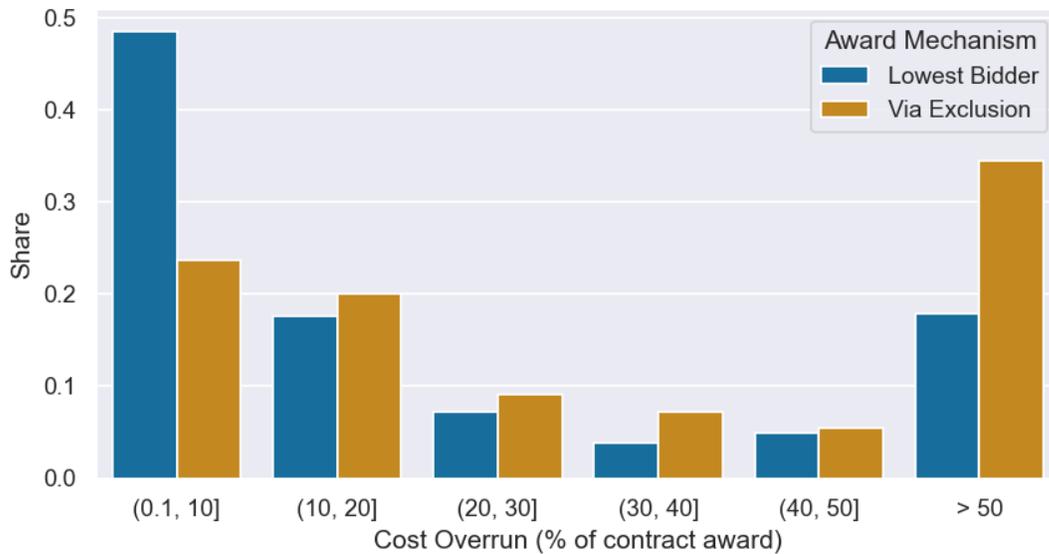
Notes — This figure plots the relationship between a contract payment outcome and the experience level of the winning firm. The y-axis is the deviation of the amount paid on the contract relative to the amount awarded — the winning firm's bid — as a share of the amount awarded. The x-axis is the experience level of the winning firm. The left and right panels use the two different definitions of firm experience. See [Section 2.2.4](#) for details.

Figure 7: Payment Outcomes for Contracts Awarded via Exclusion



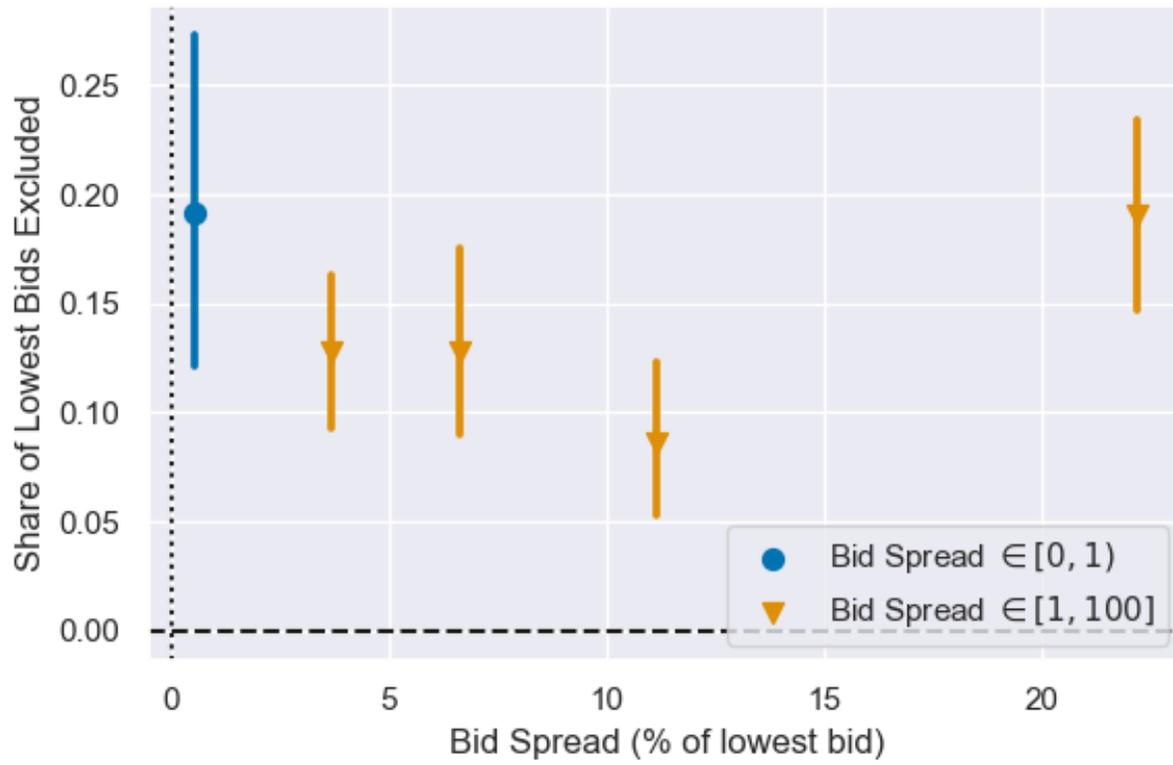
Notes — This figure shows the relationship between a contract’s payment outcome and the mechanism by which the contract is awarded. The contracts are either awarded to the lowest bid – so no exclusion – or awarded to a higher bidder that was not excluded – via exclusion. The reported p -value is from a Chi-squared test of independence for the 2-by-3 frequency table of contract counts. See [Section 2.3](#) for details.

Figure 8: Payment Overrun Magnitudes, by Award Mechanism



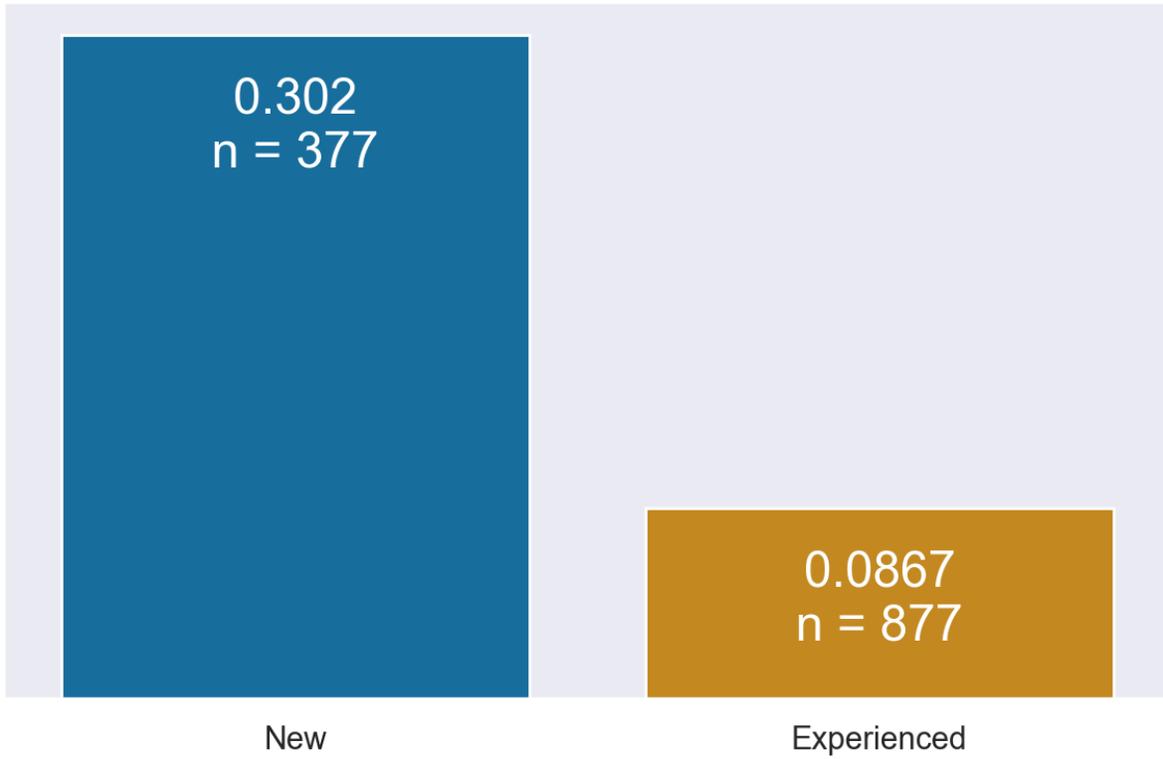
Notes — This figure shows the conditional distribution of contract payment outcomes by the mechanism in which the contract is awarded conditional on having a payment overrun. A payment overrun is define as a payment being more than 0.1% greater than the contract award amount. The magnitudes of the payment overrun are binned in 10 percentage point bins and all payment overruns over 50%. See [Section 2.3](#) for details.

Figure 9: Rate of Bid Exclusion and the Bid Spread



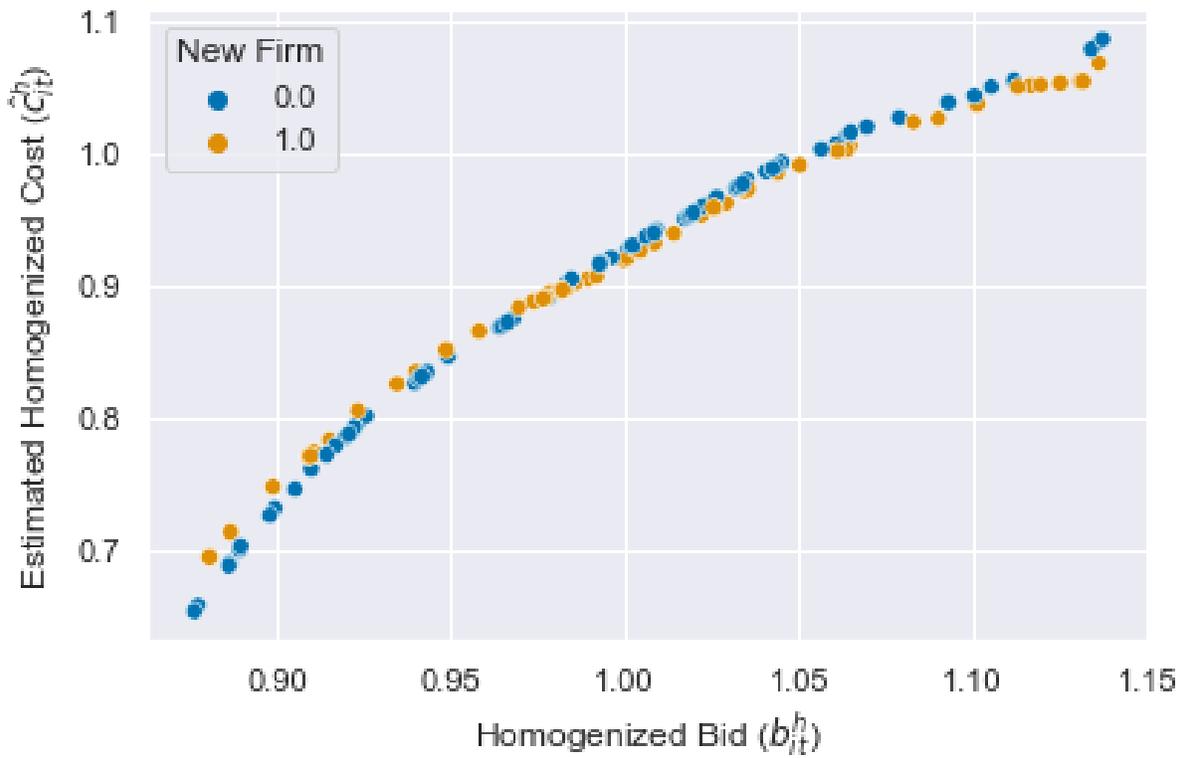
Notes — See [Section 2.4](#) for details.

Figure 10: Estimated Probability of Bid Exclusion



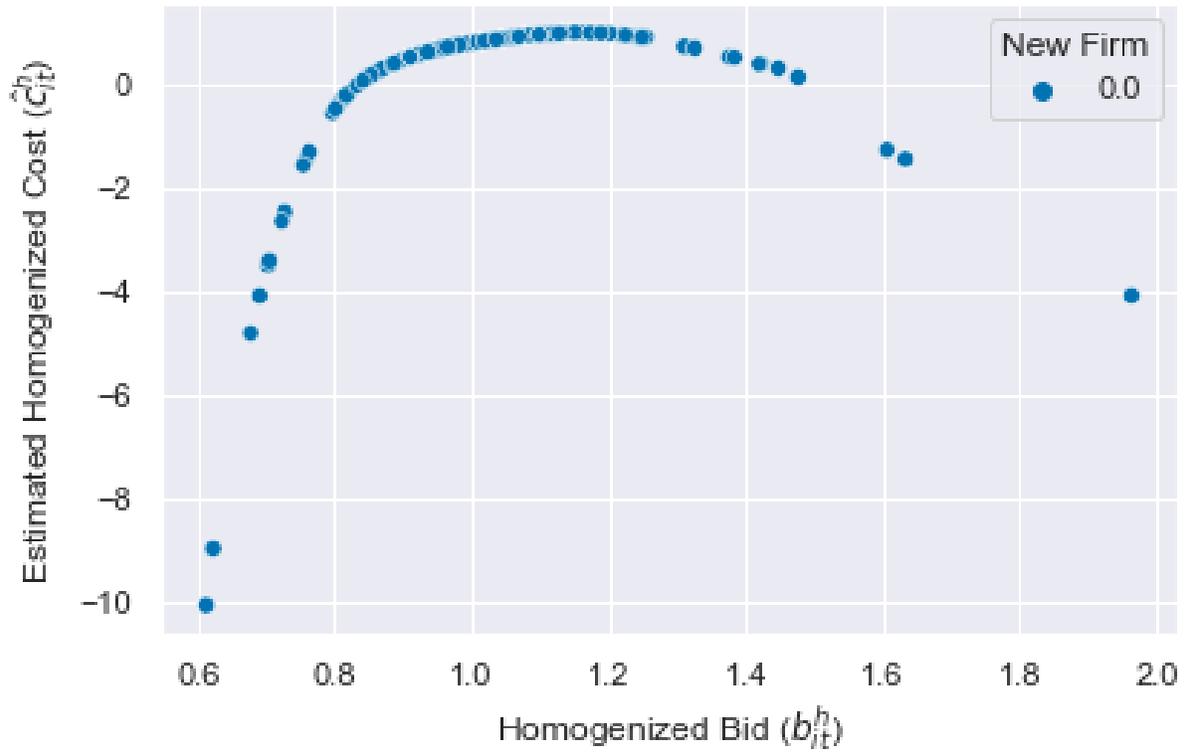
Notes — This figure shows the two firm types used in the model estimation – new firms vs. experienced firms and their estimated rates of exclusion, \hat{p}_1 and \hat{p}_2 . See [Section 4.2.1](#) for details.

Figure 11: Identified Homogenized Pseudo-Costs



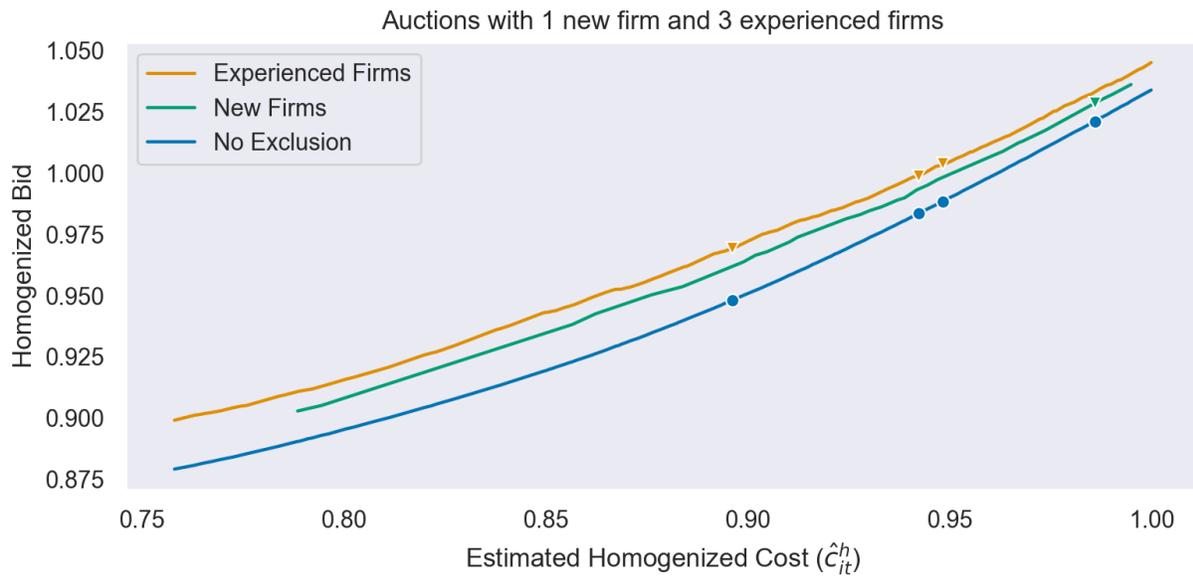
Notes — See [Section 4.2.4](#) for details.

Figure 12: Un-Identified Homogenized Pseudo-Costs



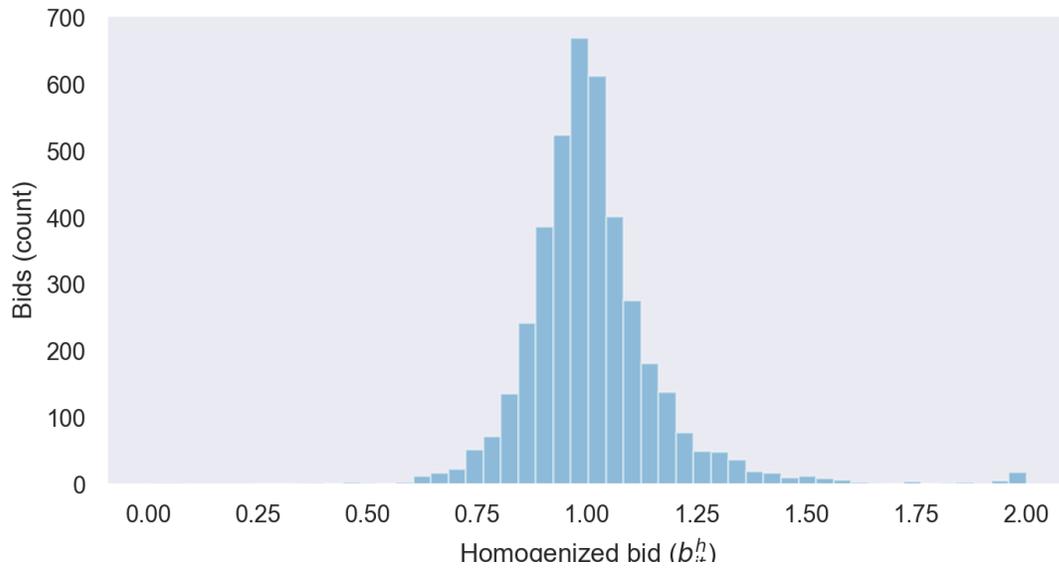
Notes — See [Section 4.2.4](#) for details.

Figure 13: Estimated Bidding Strategies Under the Two Regimes



Notes — See [Section 5.3](#) for details.

Figure 14: Histogram of Homogenized Bids b_{it}^h



Notes — Homogenized bids are clipped above $b_{it}^h = 2$. See [Section 4.1.2](#) for details.

Tables

Table 1: Bid Exclusion Example

Firm	Bid	Auction Outcome
1	\$1,546,742.50	Bid Excluded
2	\$1,656,150.00	Awarded Contract No. 15432
3	\$1,982,450.00	
4	\$2,633,775.00	
5	\$2,784,475.00	

Notes — This table demonstrates how bid exclusion is identified in the data. The bid tabulation data records the firms that submitted bids in each auction and their bid amount. This auction is assigned Specification No. 55614 (Air Filters and Accessories). In the contracts data, each contract has a specification number which allows it to be matched to an auction. The firm that is identified in the contract data is the winning bidder. In this example, Contract No. 15432 was awarded to Firm 2 for \$1.65 million even though it was underbid by Firm 1. Firm 1 is therefore identified as having had their bid excluded. The direct cost of this exclusion is roughly \$109,000. See [Section 1.2.3](#) for details.

Table 2: Bid Summary Statistics

Panel A — All Bids

	Count	Min	Max	50%	Mean	Std
Amount Bid (\$)	4,898	132.990	189,484,672.560	1,947,000.500	5,613,598.291	13,769,832.561
Chicago Bid	4,660	0	1	1	0.549	0.498
New Firm Bid	4,898	0	1	0	0.320	0.466
Excluded Bid	4,898	0	1	0	0.050	0.218
No. Prior Auctions	4,898	0	337	12	33.323	48.059
No. Prior Contracts	4,898	0	176	3	10.287	19.363
No. Firms in Auction	4,898	2	11	4	4.795	2.115
No. Chicago Firms In Auction	4,898	0	9	2	2.589	1.920
No. New Firms In Auction	4,898	0	11	1	1.480	1.601

Panel B — Lowest Bids

	Count	Min	Max	50%	Mean	Std
Amount Bid (\$)	1,254	132.990	149,916,844	1,449,444.700	4,911,373.422	12,540,642.468
Chicago Bid	1,212	0	1	1	0.528	0.499
New Firm Bid	1,254	0	1	0	0.301	0.459
Excluded Bid	1,254	0	1	0	0.152	0.359
No. Prior Auctions	1,254	0	335	9	30.036	48.618
No. Prior Contracts	1,254	0	174	3	10.894	20.719
No. Firms in Auction	1,254	2	11	3	3.906	1.864
No. Chicago Firms In Auction	1,254	0	9	2	2.041	1.735
No. New Firms In Auction	1,254	0	11	1	1.249	1.345

Notes — This table presents summary statistics for the bids in the analysis sample of auctions. Panel A presents statistics estimated on all bids in the sample. Panel B presents statistics on the subset of lowest bids, one for each auction. Under a standard reverse auction, these bids would be awarded the contract. See [Section 1.2.5](#) for details.

Table 3: Direct Costs of Bid Exclusion

	Full		Restricted	
	\$	%	\$	%
Total	57,662,593.690		38,168,461.460	
Mean	303,487.335	341.322	229,930.491	20.283
Std	768,233.665	2,143.901	611,578.626	21.167
10th	2,058.800	1.286	2,058.800	1.286
25th	7,149	4.405	5,757.750	3.772
50th	46,919.310	15.868	34,445	11.483
75th	230,782.890	44.602	185,235.888	32.996
90th	599,186	147.756	599,186	147.756
Count	190	190	166	166

Notes — This tables presents summary statistics of the direct costs of bid exclusion as observed in the analysis sample. The first column displays the summary statistics for all auctions awarded via exclusion. The second column restricts the distribution to only auctions where the lowest bid is "close" to the second lowest bid, dropping all lowest bids that are less than half the second lowest bid. See [Section 2.1](#) for details.

Table 4: Total Cost Decomposition

Case	Observed Bids	(Direct Cost) _t	(Indirect Cost) _t
No Exclusion, Same Winner	$b_{it}^h = b_{jt}^h$	0	$b_{jt}^h - \hat{\beta}_{jt}^{h, \text{IPV}} > 0$
No Exclusion, Winner Changes	$b_{it}^h < b_{jt}^h$	0	$b_{it}^h - \hat{\beta}_{jt}^{h, \text{IPV}} > 0$
Exclusion	$b_{it}^h > b_{jt}^h$	$b_{it}^h - b_{jt}^h > 0$	$b_{jt}^h - \hat{\beta}_{jt}^{h, \text{IPV}} > 0$

Notes — This table shows how the total cost of bid exclusion is divided between the direct cost, which is incurred when the procuring entity excludes the lowest observed bid in favor of a higher bid, and the indirect cost, which is the unobserved difference between the lowest bidding firm in the auction and the lowest bidding firm in the counterfactual auction with no bid exclusion. The direct cost is only positive when the procuring entity intervenes, but the indirect cost is always positive. See [Section 5.2.1](#) for details.

Table 5: Indirect Costs of Bid Exclusion

	\$	%
Total	69,769,036.790	
Mean	198,772.180	3.662
Std	451,923.780	3.192
10th	1,778.650	1.012
25th	14,239.790	1.582
50th	64,089.060	2.354
75th	184,406.760	4.846
90th	454,533.710	7.800
Count	351	351

Notes — This table provides summary statistics for the distribution of estimated indirect costs, $(\text{Indirect Cost})_t$ for the subset of auctions where the counterfactual winning bid $\hat{\beta}_{jt}^{h, \text{IPV}}$ can be calculated. See [Section 5](#) for details.