Competition and Contract Performance: Evidence from US Defense Procurement

Rodrigo Carril* Andres Gonzalez-Lira† Michael S. Walker‡

July 18, 2020

Preliminary and Incomplete
Please do not Cite or Circulate

Abstract

We study the effects of intensifying competition for contracts through advertising. Publicizing contract opportunities promotes bidder participation, potentially leading to lower acquisition costs. Yet extensive advertising could exacerbate the adverse selection of bidders on non-contractible quality dimensions. We study this trade-off in the context of procurement contracts for the U.S. Department of Defense. Our empirical strategy leverages a regulation that mandates agencies to publicize contract opportunities that exceed a certain threshold. We find that publicized contract opportunities increase competition and leads to a different pool of vendors, which, on average, offer lower prices. However, we also find that the post-award performance of publicized contracts worsens, resulting in more post-award cost overruns and delays. The latter effect is driven by goods and services that are relatively more complex, highlighting the role of contract incompleteness. To complement our findings, we develop and estimate a model in which the buyer chooses the extent of competition, and the invited firms decide on auction participation and bidding. We recover sellers' cost and quality distributions, and public buyers' preference parameters over price, quality, and idiosyncratic favoritism, and further study the extent to which the trade-off between competition and adverse selection can be delegated to the agent.

*Department of Economics, Stanford University. Email: rcarril@stanford.edu
†Haas School of Business, UC Berkeley. Email: andres.gonzalez@berkeley.edu
‡US Department of Defense. Email: michael.s.walker110.mil@mail.mil

1. Introduction

Asymmetric information is a prominent feature of many market transactions. When procuring a good or service, a buyer often deals with two sources of uncertainty: *ex-ante* sellers' production costs, and the level of (non-contractible) quality experienced *ex-post* (Laffont and Tirole, 1990; Hart and Moore, 1988). Promoting competition between potential sellers is a common way of reducing *ex-ante* adverse selection for the buyer, as evidenced by the pervasive use of bidding procedures in public procurement. But how does competition affect the equilibrium provision of non-contractible quality? If low-cost suppliers also deliver superior quality, policies that promote competition in contracting will lead to cost savings *and* improved performance. Yet if this relationship operates in the opposite way, stronger competition may "backfire", leading to the selection of sellers that perform poorly *ex-post*.

We study the interplay of competition and contract outcomes in the context of US Department of Defense (DOD) procurement. We develop a framework for nonparametrically identifying the effect of a regulation that requires agencies to publicize contract opportunities that are expected to exceed $25,000 through a centralized online platform. We exploit the discontinuous nature of these publicity requirements to estimate the effect of enhanced information diffusion about contract opportunities on four sets of outcomes: the level of competition for the award, characteristics of the buyer-contractor relationship, procurement costs, and post-award contractor performance. By providing evidence on all of these fronts, we characterize the trade-off involved in broadly advertising contract opportunities with the goal of increasing competition. Furthermore, we exploit the rich heterogeneity in the types of contracts that the DOD awards to assess the role of contract incompleteness in explaining our results.

The first part of our analysis proposes a framework for identifying the effect of publicity requirements on equilibrium award prices by investigating the observed contract price densities of publicized and unpublicized contracts.[1]The second part concerns estimating the effects of publicizing contract opportunities on three sets of outcomes: the level of competition, the characteristics of the selected vendors and their relationship to the specific buyers, and post-award performance. We implement RDD on these non-price outcomes, and argue that we can use the results of the first part of the analysis to adjust our estimates, accounting for price effects and small strategic sorting.

Implementing the proposed method, we find evidence that publicizing solicitations leads to lower contract prices. These effects are moderate, nonetheless, with an average saving of $470 at the threshold (roughly 2% of the total award). We show that contract awards advertised in the government platform attract significantly more bids, therefore achieving the desired goal of increasing competition. We also show that these marginal participants are competitive, leading to changes in the characteristics of winning firms. In particular, awardees of publicized solicitations

[1]Our method is robust to the existence of endogenous sorting below the threshold aimed to avoid publicizing certain contracts. In fact, we separately quantify the extent of strategic sorting and the price effects of publicizing contracts.

are more likely to have fewer past contracts with the awarding office and are more distant in terms of geographical location. Finally, advertised contracts result in worse ex-post performance, experiencing higher levels of cost overruns and delays. The latter results are driven by service contracts—as opposed to good purchases—, and by contracts that we ex-ante characterize as more complex. These heterogeneities highlight the role of contract incompleteness in explaining the effects of competition on post-award performance.

Taken together, our results indeed suggest that competition has mixed effects on contract outcomes. Ex-ante cost efficiencies are generated at the cost of an ex-post decrease in delivered quality. However, we show quantitatively that, for our setting, the former benefits are small relative to the latter costs. We illustrate this by comparing ex-ante price reductions with ex-post increases in cost overruns, which we can add up to compute net changes in acquisition costs. While there is a trade-off, in the sense that larger price reductions are associated with larger cost overruns, we find a net increase in total procurement costs for most product categories in our sample.

Moreover, given the volume of contracts impacted by this regulation, its implications are meaningful from a policy perspective.[2] As a reference, in 2018, the DOD publicized in the online platform contract solicitations valued in $ 5.56 billion dollars.[3] Simple calculations suggest that the price reductions due to enhanced competition as a result of publicity regulation generate annual cost savings in the order of $27.4 million. The increased overruns represent roughly $138.8 million per year. As a result, publicity requirements regulation, aimed at enhancing competition and broaden industry participation, unintendedly causes $111.4 million of annual overspending of taxpayer's money.

Our findings imply that there's an asymmetry between the incumbent sellers and the ones that bid once the solicitation is openly publicized. Furthermore, at least in some cases, it might be convenient to restrict procurement competition to a limited number of "known" bidders. Although, the extent to which this is desirable depends on the buyers' specific preferences. In future work, we will develop and estimate an equilibrium model to further understand the micro-foundations of these data patterns, and gauge the implications of policy counterfactuals. Our model will consist of three stages: First, the buyer selects the extent of competition based on it's expected outcomes. Second, heterogeneous firms decide whether to participate and bid in the auction. Finally, the performance of the selected contractor is revealed. With the estimated parameters, we can empirically quantify the factors determining the extent of competition observed in the data. In particular, we will recover sellers' cost and quality distributions, and public buyers' preference parameters over price, quality, and idiosyncratic favoritism, and further study how competition is endogenously promoted, and the extent to which the trade-off between competition and adverse selection can be delegated to the agent. Most of the details and results of the proposed model are omitted in this draft as it is still work in progress, but we expect to discuss them in upcoming

[2] The price range included our sample covers 65% of non-R&D, stand-alone contracts signed by the Army, the Navy, and the Air Force.

[3] This figure only includes stand-alone non-R&D contracts contract values between $10,000 and $1 million signed by the Army, Navy, Air Force, and the Defense Logistic Agency.

versions.

Our paper contributes to the literature examining transactions under incomplete contracting (Williamson, 1976; Hart and Moore, 1988). The existing research emphasizes that when the nature of the transaction is complex, the concerns regarding the ex-post implementation outweigh the benefits in terms of price reductions ex-ante, in which case alternative awarding mechanisms (e.g., negotiations) would be preferable (Spulber, 1990; Bajari, McMillan, and Tadelis, 2009). Nevertheless, the product and services that govern most transactions are not complex enough to depart from competitive procedures. The latter explains the prevalent use of competitive bidding by private and public organizations, even for products and services that are customized. Unlike the existing literature, our research keeps the awarding mechanism constant but analyzes the implications of varying the degree of competition in this "middle ground" that covers the vast majority of transactions. Furthermore, we study the consequences of enhancing competition across a rich set of heterogeneous goods and services, explicitly showing how the benefits and costs of competition depend on the degree of complexity of the purchase.

We also contribute to a series of studies that assess the effects of competition on prices, studying policies and rules oriented to increase (or restrict) competition in procurement (Athey, Coey, and Levin, 2013; Li and Zheng, 2009, 2012; Krasnokutskaya and Seim, 2011). We leverage variation in the intensity of competition that emerges from exogenous changes in the extent of information dissemination about contract opportunities, keeping fixed other characteristics of the institutional setting. In this respect, our paper is closely related to Coviello and Mariniello (2014), who study a similar policy in Italy.

The existing literature has also examined the role of buyers' characteristics and choices in procurement outcomes (Bandiera, Prat, and Valletti, 2009; Liebman and Mahoney, 2017; Coviello and Gagliarducci, 2017; Best, Hjort, and Szakonyi, 2017; Decarolis, Giuffrida, Iossa, Mollisi, and Spagnolo, 2018; Carril, 2020). By analyzing how public buyers endogenously promote competition, our contribution is closely related to Kang and Miller (2017), who study buyers' competition determination for IT contracts in the United States.

Finally, part of our empirical framework builds upon existing papers on nonparametric identification of behavioral parameters using density analysis (Saez, 2010; Kleven and Waseem, 2013; Kleven, 2016). Our methods contribute to the existing papers by proposing an empirical framework that recovers both contract manipulation (bunching) and the treatment effects on the assignment variable. To our knowledge, the latter is is represents a contribution to the existing literature.

The rest of the paper proceeds as follows. Section 2 provides some background on the US procurement system and the data used in our analysis. In Section 3, we present our empirical framework for estimating the effects of publicizing contract opportunities. We present our results in Section 4. Section 4.3 contrasts the pre-award price effects with the post-award performance effects. Finally, Section 5 concludes, highlighting open questions for future work.

2. Setting and Data

2.1. US Federal Procurement and Publicizing Requirements

Procurement is a large component of the US federal budget. In fiscal year 2018, federal contract awards totaled $835 billion, representing 20% of total federal outlays, and two-thirds of discretionary spending.[4] Contracts are awarded at highly decentralized levels, with more than 3,000 different contracting offices that are part of an executive or independent agency.[5] The workforce in charge of public contracting is made up of over 35,000 contracting officers whose primary role is to plan, carry out, and follow-up on purchases made by their units.

Contracting officers' scope of action is defined and limited by the Federal Acquisition Regulation (FAR). The FAR lays out policy goals and guiding principles, as well as a uniform set of detailed policies and procedures to guide the procurement process. Our analysis leverages a specific section of the FAR—Part 5 (*Publicizing Contract Actions*)—as a convenient source of quasi-experimental variation to study the effect of information diffusion.

FAR Part 5 requires publicizing contract opportunities in order to "increase competition", "broaden industry participation", and "assist small businesses (and other minority businesses) in obtaining contracts". Since October 1, 2001, contract actions that exceed $25,000 must be synopsized in an online government-wide platform which we will refer to as FedBizOpps (or FBO).[6] This implies uploading a request for quotes with a full description of the good or service being requested, and the instructions to submit the bids. We will refer to this synopsis document as a contract *solicitation*. Most of the contracts in this range are awarded to the lowest price quote that is technically acceptable according to the specifications.

Contracts that are not expected to exceed this threshold need not be publicized in FedBizOpps, although procurement officers should advertise the solicitation "by displaying [it] in a public place." This includes, for example, a physical bulletin board located at the contracting office. Of course, officers with contracts expected to fall below the threshold are still free to use FedBizOpps. On the other hand, the regulation allows for exemptions to the requirement above the threshold, as long as the procurement officer can properly justify it on grounds that it "compromises national security", that "the nature of the file does not make it cost-effective or practicable", or that "it is not in the government's interest". Therefore, while this policy discretely affects the likelihood of publicized contracts around the threshold, we anticipate that compliance may be far from perfect, given the voluntary nature of the rule below this value and the availability of exceptions above.

[4] Discretionary spending excludes mandatory programs (such as Social Security and Medicare) and interest on debt.

[5] Executive agencies are headed by a Cabinet secretary, like the Department of Defense, the Department of State, or the Department of Health and Human Services. Independent agencies, which are not part of the Cabinet, include the Central Intelligence Agency, the Environmental Protection Agency, and the Federal Trade Commission.

[6] Throughout our period of analysis, this online platform—designated as the "government point of entry" by the FAR—was called Federal Business Opportunitites (FBO) and was available at: `fedbizopps.gov`. In late 2019 (after our sample period ends), the government point of entry migrated to `beta.sam.gov`, featuring significant changes to the user interface.

2.2. Data

We use two complementary sources of data. The first one consists of the historical files from FedBizOpps, which provides detailed information on pre-award notices (i.e. solicitations) posted on the platform. The second one is the Federal Procurement Data System - Next Generation (FPDS-NG), which tracks federal contracts from the time of their award and including all follow-on actions, such as modifications, terminations, renewals, or exercises of options.

We merge awards from FPDS-NG to notices on FedBizOpps using the solicitation number. Note, however, that while FPDS-NG contains the universe of federal awards, FedBizOpps only has the notices posted on the website. From this matching process, we construct a dummy variable that is equal to 1 if we are able to merge a contract with any pre-award notice on FedBizOpps, in which case we say the contract was *publicized*. Appendix Figure A1 describes the typical timeline of events surrounding the life-cycle of a contract, and the appropriate data source that records that information.

In addition, we observe detailed information for each contract award, including the dollar value of the funds obligated, a four-digit code describing the product or service, codes for the agency, sub-agency, and contracting office making the purchase, the identity of the private vendor, the type of contract pricing, the extent of competition in the award, characteristics of the solicitation procedure, the number of offers received, and the applicability of a variety of laws and statutes.

The analysis sample consists of all definitive contracts[7] with award values between $ 5,000 and $ 45,000, awarded in fiscal years 2014 through 2019 by the Department of Defense (DOD),[8] for products and services other than Research and Development (R&D).[9] Table 1 presents summary statistics of the sample. In total, there are roughly 240,000 contracts awarded by 760 contracting offices to almost 60 thousand distinct firms. Contract durations are expected to be 55 days on average and are awarded on a fixed-price basis. A noteworthy feature of this setting is that competition is very limited: more than a third of the awards are set-aside for a particular type of firm (typically, small business), and the average contract receives 2.4 offers, with the median contract receiving a single offer. The Department of the Navy and the Army each account for more than 40% of the contracts, with the rest being mostly awarded by the Air Force. Winning vendors are often geographically close to the contracting offices, with both located in the same state in 3 out of every 4 contracts. Finally, 62% of suppliers are characterized as small businesses.

We also observe rich information about the type of good and service that is contracted upon. Each award is classified into one of 1,918 possible standardized 4-digit alphanumeric codes. These can be aggregated into 101 broader 2-digit product categories, 77 goods and 24 services. Table 2

[7] Federal contracts can be broadly categorized into two types: definitive contracts (DCs) and indefinite delivery vehicles (IDVs). DCs are stand-alone one-time agreements with a single vendor for the purchase of goods or services under specified terms and conditions. See Carril (2020) for more details. We simplify the analysis by focusing exclusively on DCs, which are well-defined requirements involving a bilateral relationship within a single government unit and a private firm.

[8] The Department of Defense represents 55% overall federal spending and more than 60% in the restricted sample.

[9] R&D awards are subject to a whole set of special rules, see FAR Part 35.

shows the top 10 most common 2-digit good and service categories. The most common product categories are ADP Equipment Software, Medical Equipment and Supplies and Maintenance and Repair Equipment.

Table 1: Summary statistics

	Mean
Contract Characteristics	
Expected Award Amount	22,070
Expected Duration (days)	55.15
Fixed-Price Contract	0.999
Competitively Awarded	0.614
Set Aside Award	0.357
Simplified Procedure	0.728
Competition	
Number of Offers	2.452
One Offer	0.530
Contracting Office Characteristics	
Navy	0.422
Army	0.402
Air Force	0.134
Other	0.043
Awarded Firm Characteristics	
Foreign	0.092
Within-State Firm	0.741
Small Business	0.620
Woman Owned Business	0.137
Sample	
No. of Contracts	240,514
No. of Contracting Offices	760
No. of Firms	59,697

Notes: This table presents summary statistics. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2014 through 2019. An observation is a contract, defined by aggregating all contract *actions* (initial award, modification, termination, etc.) associated with the same contract ID.

Table 2: Top product and service categories

	Goods		Services	
Rank	Name	N Contracts/year	Name	N Contracts/year
1	ADP Equipment and Software	3,005	Maintenance/Repair of Equipment	2,430
2	Medical Equipment and Supplies	2,998	Support Services (Professional)	1,187
3	Laboratory Equipment	1,643	Utilities And Housekeeping	1,096
4	Electrical Equipment Compontents	1,593	Transport, Travel, Relocation	854
5	Communication/Coherent Radiation	1,202	ADP and Telecommunications	806
6	Furniture	810	Lease/Rent Equipment	753
7	Power Distribution Equipment	697	Maintenance of Real Property	688
8	Ship And Marine Equipment	574	Education And Training	560
9	Hardware And Abrasives	530	Construct Of Structures/Facilities	335
10	Construction And Building Material	459	Social Services	286

Notes: This table presents average annual counts of contracts in the most common product categories. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. An observation is a contract, defined by aggregating all contract *actions* (initial award, modification, termination, etc.) associated with the same contract ID. A 4-digit alphanumeric code (PSC) is observed for each contract. The categories listed are constructed by aggregating PSC codes to two-digits for goods, and to a single digit (letter) for services.

3. Empirical Framework

In this section, we propose a framework that captures some of the key aspects of the DOD procurement process, particularly in relation to publicizing decisions. We then use this framework to motivate and explain our empirical approach to estimating the effects of publicly soliciting contracts on awarding prices and other contract outcomes.

3.1. Setup

A public buyer seeks to award a contract k to acquire a good or service. There is a set of potential suppliers $\mathcal{J}$. Conditional on awarding contract k to supplier j, then the buyer obtains contract outcomes $Y_{jk} = (p_{jk}, q_{jk}, \psi_j)$, where p_{jk} is the award price, q_{jk} is the *realized* post-award performance, and ψ_j captures other contractor characteristics that are independent of the particular contract k.

The buyer has a set of preferred or default suppliers $\mathcal{J}^0 \subset \mathcal{J}$, for which she knows $(c_{jk}, \tilde{q}_{jk}, \psi_j)$. Alternatively, she can decide to post the solicitation on a public platform (i.e. FedBizOpps), in which case the set of potential bidders (weakly) expands to $\mathcal{J}^1 \supseteq \mathcal{J}^0$. We denote by D_k the decision to publicize a solicitation, so that D_k equals one when the solicitation is posted online, and zero otherwise. Conditional on the set of participants, the contract is awarded to the lowest bidder through a sealed bid first-price auction.

Before deciding on whether to publicize the solicitation, the buyer privately observes an independent price estimate $\tilde{p}_k$. There is a regulation that encourages the buyer to publicize the solicitation whenever the estimate exceeds a certain threshold $\bar{p}$. If $\tilde{p}_k \leq \bar{p}$, the buyer is free to choose $D_k \in \{0, 1\}$. If $\tilde{p}_k > \bar{p}$, not publicizing involves a utility cost of Δ.

Finally, our setting allows that the buyer can modify the characteristics of the purchase in order to lower the initial estimate just enough to avoid the regulation. For example, faced with an estimate that exceeds the threshold, the buyer could decide to reduce the quantity purchased so that a revised price estimate falls below $\bar{p}$. Doing so generates a utility cost of $\lambda(\tilde{p}_k)$ that is increasing in the original price estimate. We call this decision "bunching" and denote it as an indicator B_k.

All of the above implies that faced with an estimate $\tilde{p}_k$, the buyer has three possible courses of action. First, she can go ahead with the purchase "as is" without publicizing the solicitation $(D_k = 0, B_k = 0)$. Denoting by $Y^0 \equiv (p_k^0, q_k^0, \psi^0)$ the outcomes that the buyer would get in this case, the utility obtained is given by $U_k^0 \equiv U(Y_k^0) - \Delta \cdot \mathbf{1}[\tilde{p}_k > \bar{p}]$. Second, she may choose to publicly solicit the contract $(D_k = 1, B_k = 0)$, leading to outcomes that we denote $Y_k^1 \equiv (p_k^1, q_k^1, \psi^1)$, and a utility of $U_k^1 \equiv U(Y_k^1)$. A third possibility is to solicit among the preferred local contractors by bunching $(D_k = 0, B_k = 1)$, leading to outcomes that we denote $Y_k^B \equiv (p_k^B, q_k^0, \psi^0)$ and a utility of $U_k^B \equiv U(Y_k^B) - \lambda(\tilde{p}_k)$. Note that this action may only be chosen when $\tilde{p}_k > \bar{p}$.

The fundamental challenge of our empirical setting is that we do not observe price estimates $\tilde{p}_k$, nor the different prices that the buyers would obtain if they followed alternative courses of action. In other words, we only observe equilibrium transaction prices p_k, and their decision to publicize the solicitation D_k. Dropping the subscript k from all the variables, we can write the observed awards as a function of other observed and latent variables:

$$\begin{aligned} p &= D \cdot p^1(\tilde{p}) + (1 - D) \cdot \Big[(1 - B) \cdot p^0(\tilde{p}) + B \cdot p^B(\tilde{p})\Big] \\ &= p^0(\tilde{p}) + D \cdot \Big[p^1(\tilde{p}) - p^0(\tilde{p})\Big] + B \cdot (1 - D) \cdot \Big[p^B(\tilde{p}) - p^0(\tilde{p})\Big] \end{aligned}$$

Our first goal is to propose a framework for nonparametrically characterizing the distribution of $\tilde{p}$ and the price effects of publicity $(p^1(\tilde{p}) - p^0(\tilde{p}))$, from the distribution of observed awards p and publicizing decisions D. Being able to recover information about these unobserved price estimates $\tilde{p}$, we then use the discontinuous nature of the publicity regulation to assess effects on other contract outcomes of interest.

3.2. Recovering Expected Awards and Estimating Price Effects

We build upon the proposed conceptual framework to identify relevant parameters using excess and missing masses in empirical density distributions around the policy threshold. The proposed method hinges on the comparison between the empirical distribution and an estimated counterfactual distribution, using a procedure we now describe. We proceed in three steps; first, we lay out the assumptions we build upon, and then, conditional on these assumptions, we characterize the distribution of densities around the regulation threshold. Finally, we propose an estimation strategy that leverages the empirical distribution of contract values.

3.2.1. Baseline Assumptions: Our proposed method relies on the following assumptions.

A1 Price estimates are drawn from a smooth distribution, $\tilde{p} \sim F(\cdot)$ with density $f(\cdot)$.

A2 Price estimates equal the price that the buyer would obtain from preferred vendors absent any bunching, i.e. $p^0(\tilde{p}) = \tilde{p}$. In other words, $\tilde{p}$ represents an accurate estimate of the price of soliciting locally.

A3 Publicizing solicitations leads, on expectation, to a constant proportional price reduction, relative to the price of soliciting locally. In other words, $p^1(\tilde{p}) = (1 - \gamma(\tilde{p})) \cdot p^0(\tilde{p})$, with $\mathbb{E}[\gamma(\tilde{p})|\tilde{p}] = \bar{\gamma} > 0$ for all $\tilde{p}$.

A4 $\lambda(\tilde{p})$ is such that there exist $p_H > \bar{p}$ such that bunching is never chosen for $\tilde{p} > p_H$.

The first assumption is standard. It requires that the key latent density that we seek to recover be smooth. Importantly, we impose no additional parametric restrictions on $f(\cdot)$. Our second assumption is substantially stronger, requiring that buyers anticipate the award that they would obtain from a non-publicized solicitation.[10] In principle, we conjecture that this is stronger than what we actually need, which is that buyers' estimates are right on average, i.e. $\mathbb{E}\left[p^0(\tilde{p})|\tilde{p}\right] = \tilde{p}$. We are currently updating our method to make it valid under this weaker assumption.

Assumption **A3** requires that the expected price effect is constant in proportional terms. Importantly, this is true *in expectation*, so we allow different contracts with the exact same estimate to lead to very different realized awards. Because in practice our method is implemented on a relatively narrow window of award values (between $10,000 and $40,000), we think this assumption is not too restrictive. Furthermore, we are working on an extension allowing other functional forms for the price effect.

Finally, the fourth assumption considers that the cost of adjusting estimates is such that bunching will be somewhat local to the threshold. For a sufficiently large value of p_H, we can be sure that an estimate of p_H or higher will never result in bunching. This assumption is standard on the bunching literature.

3.2.2. Characterizing Award Densities: First, consider the baseline scenario in which there's no possibility to bunch, and there are no price effects of publicity. That is, suppose that $\gamma(\tilde{p}) = 0$ and that $B = 0$ for all contracts. Our assumptions imply that both p^0 and p^1 will be exactly equal to the price estimate $\tilde{p}$. This will lead to distributions of publicized and non-publicized awards such as the ones depicted in Figure 1. On the one hand, the total observed density of contract prices $h(p)$ equals the density of price estimates $f(\tilde{p})$, which is smooth. On the other hand, the regulation introduces a strict shift in the cost of not publicizing for contracts with estimates that exceed $\bar{p} = 25,000$, making the fraction of publicized solicitations to jump discontinuously at this threshold. The discrete increase in the probability of publicity translates into a discontinuous jump in the density of advertised contracts that is compensated one-to-one with a drop in the density of non-publicized contracts.

[10] The treatment promotes the publicity adoption by introducing a utility cost if the contract's value without advertising is above the threshold. Thus, the amount that determines the treatment assignment is the price of the contract without publicity.

Figure 1: No price effects or bunching

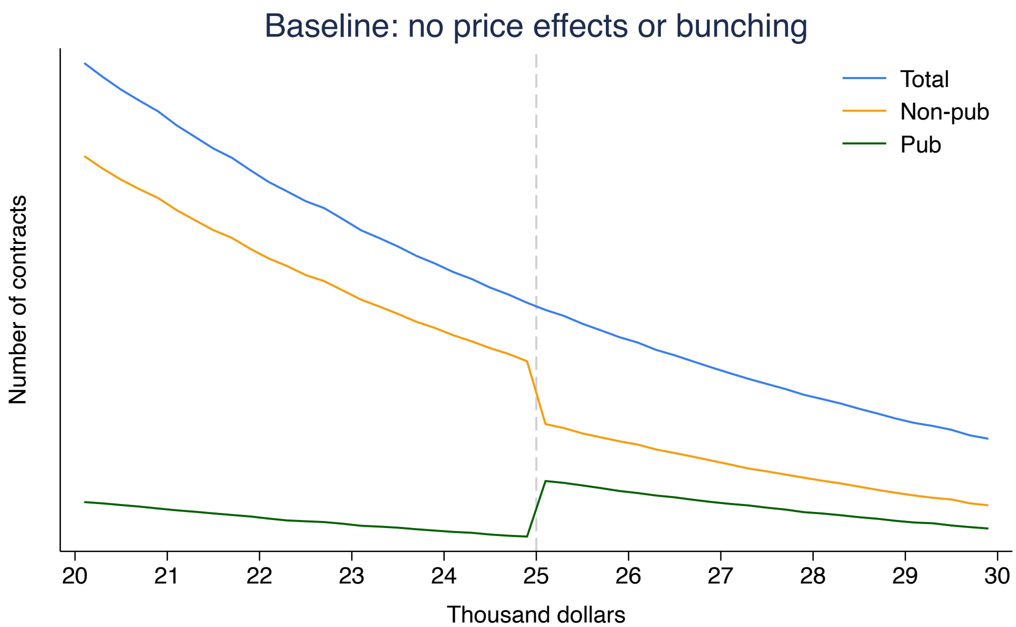

Now, suppose that we allow buyers to bunch maintaining the assumption of zero price effects. Buyers can modify the purchase characteristics to reduce the price to circumvent the regulation; this action will translate into a non-publicized award of $p^B = \bar{p}$. Figure 2 Panel (a) and panel (b) respectively show the effects of bunching on the densities of non-publicized and publicized contracts. Since bunching aims to avoid soliciting the contract publicly, it generates an excess mass of *non-publicized* awards right below the threshold, as shown in the left panel. The excess mass below the threshold in the non-publicized density has to be equal to the sum of the missing masses to the right of the threshold in the non-publicized and publicized. Finally, note that, because of **A4**, there is a point to the right of the threshold where bunching responses no longer occur, as seen by the fact that the solid and dashed lines converge in Figure 2.

Finally, consider the role of price effects on top of bunching responses. Figure 2 panel (c) and panel (d) respectively show these effects for the density of non-publicized and publicized awards. Because price effects only affect advertised contracts, the mass in panel (c) is unaffected by the size of γ. On the contrary, the full distribution of publicized awards is "shifted" by price effects. We can break down this into two margins: first, the density is proportionally shifted to the left; publicized awards are now awarded in expectation for $\mathbb{E}[p_1] = (1 - \bar{\gamma}) \cdot p_0$. Second, the sharp discontinuity is "smoothed-out", because of the variance in the individual price effects γ.

To see more clearly that the price effects of publicity can be decomposed into the two changes described, suppose that we knew the value of the average price effect, $\bar{\gamma}$. Starting from the observed density of publicized contracts that includes the price effects, we could take all contracts and adjust them proportionally by $\frac{1}{1-\bar{\gamma}}$. Figure 3 shows this exercise, and confirms that such adjustment would bring the observed distribution to the levels of the unobserved density of price estimates, except

Figure 2: Bunching and price effects

(a) Bunching for $D = 0$

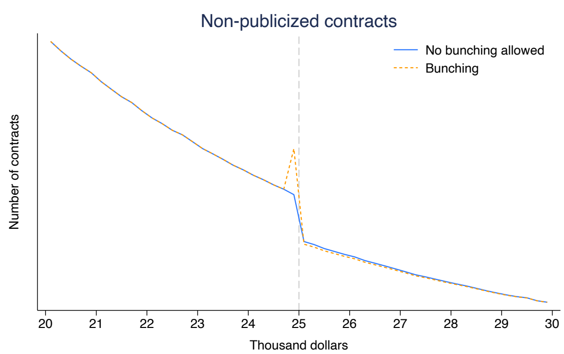

(b) Bunching for $D = 1$

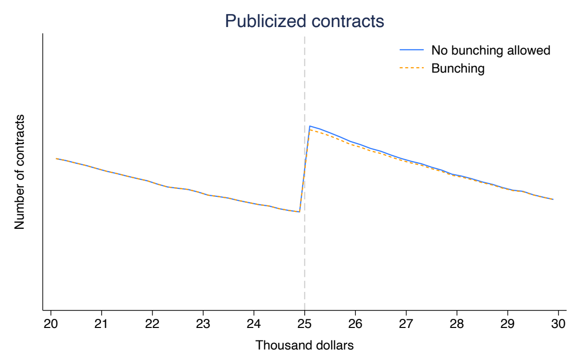

(c) Bunching + Price Effects for $D = 0$

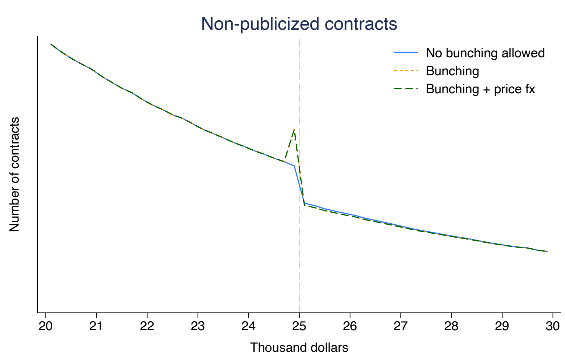

(d) Bunching + Price Effects for $D = 0$

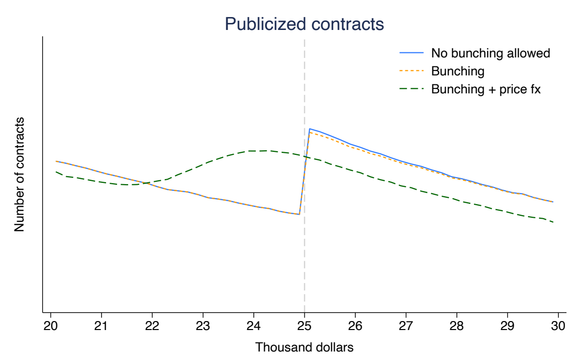

Figure 3: No price effects or bunching

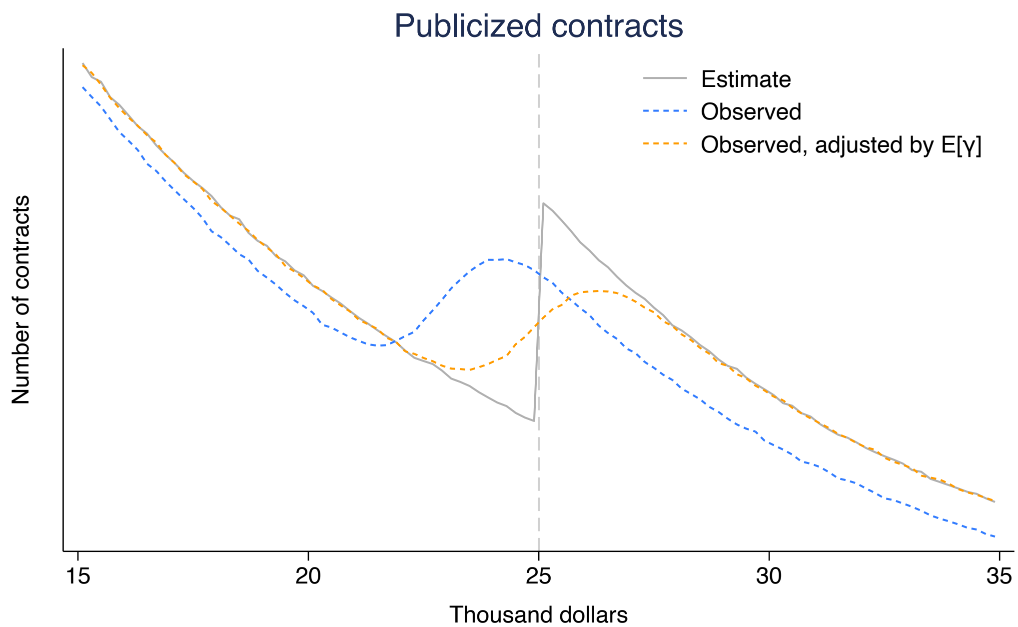

around the threshold where the discontinuity gets "smoothed-out".

In what follows, we explain our method to nonparametrically recover the density of price estimates $f(\tilde{p}|D = 0)$ and $f(\tilde{p}|D = 1)$ from the observed density of awards, $h(p|D = 0)$ and $h(p|D = 1)$. We directly build on the previous graphical analysis to relate these unobserved counterfactual magnitudes to the observed densities.

3.2.3. Estimation method: Our approach builds upon methods originally developed to measure behavioral responses to taxes (e.g. Saez, 2010; Chetty et al., 2011). In particular, it is closely related to the framework proposed by Kleven and Waseem (2013) for the case of notches. Nonetheless, our setting departs from these standard approaches in two respects. First, while standard bunching approaches seek to identify a single magnitude, namely the amount of excess bunching around the threshold,[11] we want to *separately* identify the extent of bunching from the expected price effects of publicly soliciting contracts. Intuitively, Figure 2 highlights that we can recover these two latent quantities because we observe *two* different densities that are deferentially affected by them.[12]

Second, our setting is also special because the regulatory environment generates a discontinuity in the densities, even in the absence of bunching and price effects. In other words, the latent counterfactual densities of publicized and unpublicized contracts that we seek to recover will not be smooth. The latter makes it unfeasible to use the standard interpolation methods in the traditional approaches, where the counterfactual density close to the threshold is estimated from projecting the

[11]This magnitude is relevant because under certain assumptions researches can use it to recover the value of the elasticity of taxable income.

[12]We are currently working on a formal discussion of the identification of these two objects under **A1-A4**.

observed density from both the left and the right. As we will see, we will instead have to rely on one-sided projections of the densities.

Roughly speaking, we "reverse-engineer" the logic presented in the preceding section. Starting with a guess of average price effects $\bar{\gamma}$, we transform the observed frequency distribution of publicized contracts as in Figure 3: from the frequency distribution of $(p|D = 1)$, we construct the frequency distribution of $(\frac{p}{1-\bar{\gamma}}|D = 1)$. If the guess is correct, then this adjusted distribution should coincide with the distribution of price estimates $\tilde{p}$, except near the threshold. We can then recover the counterfactual distribution around the threshold by projecting from outside some excluded areas.

On the other hand, we can do a similar projection with the distribution of non-publicized contracts, as we see in Figure 2 panel (c). Far from the threshold, the distribution of observed awards and counterfactual price estimates should coincide, and we can guess the counterfactual close to the threshold by projecting the observed distribution with a polynomial fit from each side.

Finally, two constraints must be satisfied after the estimation: First, the total number of observed awards should be equal the total number of unobserved price estimates $\tilde{p}$, and second, the size of the jump in the mass publicized contracts should exactly compensate the reduction in non-publicized contracts at the threshold. Because we started with a guess of average price estimates, our method iterates over different values of $\bar{\gamma}$ to find the one that minimizes deviations from the constraints.

More specifically, we partition the space of contracts award values between \$10,000 and \$40,000 in right-inclusive bins of \$200 width,[13] so that bin b contains contracts between $(1000 \cdot b - 200, 1000 \cdot b]$. Let n_b^D be the number of observed contracts with publicity status D in bin b. Let $\underline{b}^D$ and $\bar{b}^D$ be the lower and upper limits of the excluded regions for each distribution $D \in \{0, 1\}$, and b^T is the bin that contains the threshold, i.e., $b^T = (24, 800, 25, 000]$. The one-sided polynomial extrapolations on the frequency distribution of $D = 0$ are obtained by estimating:

$$n_b^0 = \sum_{x=0}^{Q} \alpha_x^0 \cdot b^x + \sum_{j=\underline{b}^0}^{b^T} \gamma_j^0 \cdot \mathbf{1}[b = j] + \nu_b^0, \text{ for } b = \{10.2, 10.4, ..., 25.0\} \tag{1}$$

$$n_b^0 = \sum_{x=1}^{Q} \beta_x^0 \cdot b^x + \sum_{j=b^T}^{\bar{b}^0} \delta_j^0 \cdot \mathbf{1}[b = j] + \eta_b^0, \text{ for } b = \{25.2, 25.4, ..., 40.0\} \tag{2}$$

and then computing the fitted values ignoring the contribution of the excluded area, so that the estimated frequency distribution is:[14]

$$\widehat{n}_b^0 = \begin{cases} \sum_{x=1}^{Q} \widehat{\alpha}_x^0 \cdot b^x & \text{for } b = \{10.2, 10.4, ..., 25.0\} \\ \sum_{x=1}^{Q} \widehat{\beta}_x^0 \cdot b^x & \text{for } b = \{25.2, 25.4, ..., 40.0\} \end{cases} \tag{3}$$

[13]That is, 150 price bins: $\{(10000, 10200], (10200, 10400], ...(39800, 40000]\}$.

[14]Note that regression (1) extrapolates from the left, and equation (2) to the right of the excluded area, by fitting a polynomial regression each side independently we enhance the flexibility of the model.

For the publicized contracts ($D = 1$), the process is analogous, except that we don't use the observed frequency distribution $\{n_b^1\}_{b=10}^{40}$, but rather an adjusted frequency distribution $\{m_b^1(\bar{\gamma})\}_{b=10}^{40}$, which is obtained by multiplying the award value of each publicized contract by $(\frac{1}{1-\bar{\gamma}})$. By the same procedure above we obtain fitted values $\widehat{m}_b^1(\bar{\gamma})$.

Finally, there are two constraints that we need to impose. The first one is that the total amount of contracts in the counterfactual distribution equals the observed ones:

$$c_1(\bar{\gamma}) = \sum_b (\widehat{n}_b^0 + \widehat{m}_b^1(\bar{\gamma})) - \sum_b (n_b^0 + n_b^1) = 0 \tag{4}$$

The second constraint is that the discontinuous changes in the densities at the threshold exactly compensate each other:[15]

$$c_2(\bar{\gamma}) = \left[\widehat{n}_{b^T}^0 - \widehat{n}_{b^T+1}^0\right] - \left[\widehat{m}_{b^T+1}^1(\bar{\gamma}) - \widehat{m}_{b^T}^1(\bar{\gamma})\right] = 0 \tag{5}$$

We search for values of $\bar{\gamma}$ to minimize a standard distance metric $\Omega(\bar{\gamma}) = c(\bar{\gamma})'Wc(\bar{\gamma})$, where $c(\bar{\gamma}) = (c_1(\bar{\gamma}), c_2(\bar{\gamma}))'$ and W is a weighting matrix. The parameters of this procedure are the degree of the polynomial Q and the limits of the excluded areas $(\underline{b}^D, \bar{b}^D)$. The estimated price effects are local as they stem from the set of contracts that are manipulated by the regulation at the discontinuity.

3.3. Estimating Effects on Non-price Outcomes

Through the described procedure, we can recover an estimate of the distribution of unobserved price estimates $\tilde{p}_k$. This is useful because we can leverage the discontinuous nature of the publicity requirements to gauge the effects of publicity on a set of other relevant outcomes, including the level of competition, characteristics of the winning bidder, and post-award contractor performance.

Consider specifications of the following form:

$$Y_k = \alpha + \beta \cdot D_k + g(\tilde{p}_k) + X_k'\delta + \epsilon_k \quad , \tag{6}$$

Where we are interested in β, the causal effect of publicizing a solicitation on contract outcome Y_k. In the standard Regression Discontinuity Design (RDD), we obtain an estimate of $\widehat{\beta}_{LATE}$ by instrumenting D_k with the discontinuity in publicity requirements. The first-stage of this IV procedure is of the form:

$$D_k = \lambda + \gamma \cdot \mathbf{1}[\tilde{p}_k > \bar{p}] + g(\tilde{p}_k) + X_k'\eta + \nu_k \quad , \tag{7}$$

[15] This constraint stems directly from the assumption that the underlying distribution of $\tilde{p}$ is smooth at the threshold, i.e., increases in the mass publicized contracts need to be exactly compensated by reductions of the mass of non-publicized contracts.

For some smooth function $g(\cdot)$. A key advantage of this approach is that it is possible to provide compelling evidence on the existence of an effect by graphically showing the reduced form of this model, i.e.:

$$Y_k = \mu + \phi \cdot \mathbf{1}[\tilde{p}_k > \bar{p}] + g(\tilde{p}_k) + X_k' \pi + \xi_k \quad . \tag{8}$$

Again, the key challenge we face is that we observe p_k, but not $\tilde{p}_k$. However, it is easy to see why our previous method helps us adress this empirical challenge, by allowing us to recover an unbiased estimate of $\tilde{p}_k$. For publicized contracts, we want $\mathbb{E}[\tilde{p}_k|p_k]$. By **A3**, this equals $\mathbb{E}[p_k(1-\bar{\gamma}_k)|p_k] = p_k \cdot (1 - \mathbb{E}[\bar{\gamma}_k|p_k])$, where p_k is observed and from our previous analysis we can recover a consistent estimate of $\mathbb{E}[\bar{\gamma}_k|p_k]$.

On the other hand, estimating the extent of bunching allows us to re-weight observations to appropriately reflect the average outcomes in the absence of this behavior. Using the results from Gerard, Rokkanen, and Rothe (2016), we can nonparametrically identify bounds on the treatment effects.

In this version of the paper, we do not implement these corrections, leaving them for future work. We will present in Section 4.2 the results from RDDs that use realized awards as the running variable, despite the interpretation caveats highlighted here.

However, it is important to consider that, in practice, the results presented below in Section 4.1 will show that both the price effects and the extent of bunching are small in magnitude in our sample. These are interesting results in and of themself and imply that the corrections mentioned above of the RDDs will have at most modest implications on the estimated effects. So, while the estimates presented in Section 4.2 will not be definitive, we expect them to be highly informative of the true causal effects.

4. (Preliminary) Results

4.1. Price Effects of Competition

4.1.1. Estimation Details: To implement the method described in Section 3.2.3, we choose the following values for the estimation parameters. We use a 5th-degree polynomial ($Q = 5$), an asymmetric excluded region for non-publicized awards of $(\underline{b}^0, \bar{b}^0) = (22, 30)$, and a symmetric window for publicized awards of $(\underline{b}^1, \bar{b}^1) = (20, 30)$. In future iterations we will provide a full sensitivity analysis to these parameter choices. We also incorporate round number corrections to abstract from systematic patterns commonly observed in binned frequency distributions. In particular, we follow Kleven and Waseem (2013) and Carril (2020) in working with smoothed distributions that eliminate round-number effects.

4.1.2. Results: Our main estimates are presented in Table 3, both for the full sample, and separately for goods and services. The main finding is that the average price effect of publicizing contracts is 0.019. This means that, on average, posting a solicitation on FedBizOpps reduces the awarded price by roughly 2%. At the threshold, this would mean a price reduction of $470.

Table 3: Parameter estimates

	Estimates		
	All Products	Goods	Services
Mean price effect ($\bar{\gamma}$)	0.019	0.016	0.022
	(.)	(.)	(.)
Excess bunching ($Pr(B=1)$)	0.026	0.028	0.021
	(.)	(.)	(.)

Notes: This table presents model parameter estimates obtained via simulated method of moments. The estimation sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2011 through 2017. The moments used in the estimation correspond to the smoothed fraction of publicized and unpublicized contracts, for each $500-wide award bin between $10,000 and $40,000.

Figure 4 shows our estimated counterfactual distribution of awards. Panel (a) shows the density of publicized contracts, along with a counterfactual density in which price effects are eliminated. Under our modelling assumptions, this coincides with the ex-ante price estimates observed by the buyers. As we discussed in Section 3.2, price effects smooth-out this density at the threshold. The excess mass relative to the counterfactual that we see to the right of the threshold represents the set of contracts that had expected awards above the threshold, but were brought below it because of the competitive effect of publicizing the solicitations.

Panel (b) repeats the exercise for the non-publicized contracts, this time shutting down the possibility of bunching. Again, this counterfactual reflects the price estimates $\tilde{p}$. As expected, some of the mass just below the threshold relocates to the right (the rest goes to the publicized density to satisfy the integration constraint).

Finally, we present the effect of each of the two forces (bunching and price effects) on the total density of contracts in Panel (c). Based on this analysis, we can conclude that one-third of the total excess mass below the threshold is due to price effects of competition, whereas the rest is explained by strategic bunching. Removing both effects gives us an estimated density of price estimates $\tilde{p}$ that is completely smooth.

Figure 4: Counterfactuals

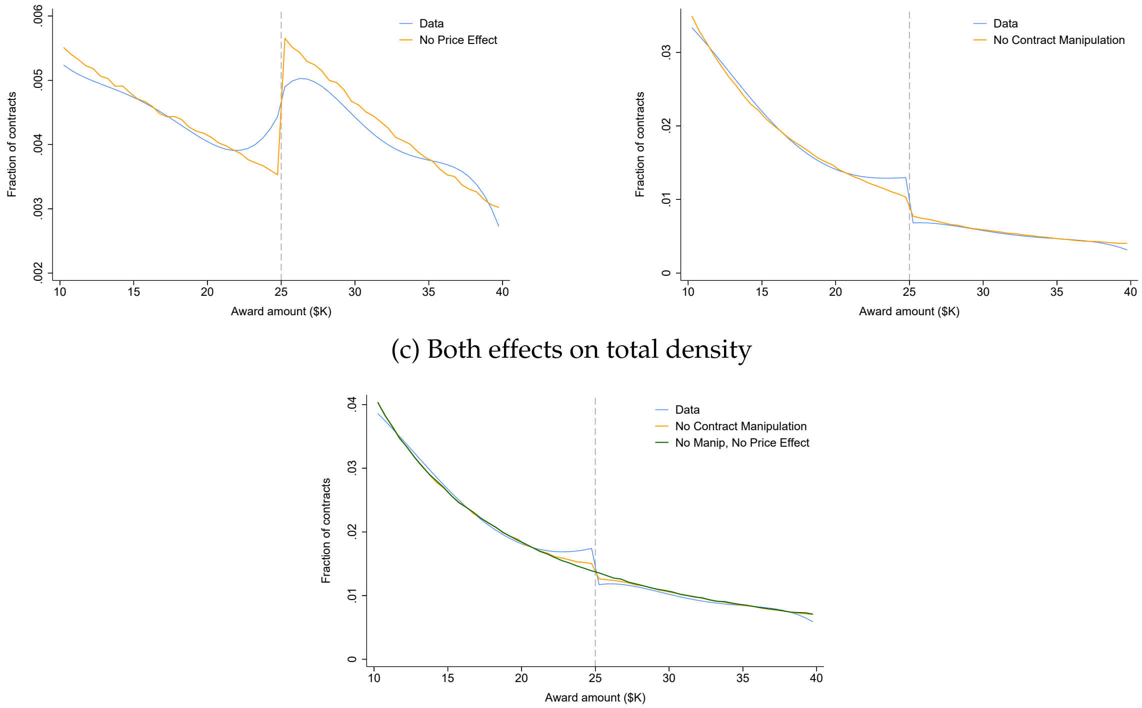

(a) Price effects on publicized contracts

(b) Bunching of unpublicized contracts

(c) Both effects on total density

Notes: This figure shows model-based counterfactual estimates that seek to quantify the effects of price effects and strategic bunching. Panel (a) shows the smoothed density of publicized contracts in the data, compared to a simulated density assuming a price effect of zero for all contracts. In Panel (b), we plot the smoothed density of unpublicized contracts in the data, compared to a simulated density assuming no contract manipulation. Panel (c) shows the effects of both of these counterfactual simulations on the total density of contracts.

4.2. Contract Publicity, Competition and Outcomes

We estimate the effect of publicizing contract solicitations on procurement outcomes following the strategy described in Section 3.3, consisting of a set of three key specifications, Equation (6), Equation (7) and Equation (8). We consider linear and quadratic fits for the function $g(\cdot)$, and do not include controls X_k in our baseline specification. We estimate these regressions using adjusted award amounts, as discussed, the price effects are quantitatively small, so the unadjusted figures are almost identical.[16]

We start by estimating the first stage Equation (7). The results are presented graphically in Figure 5. We see that the use of FedBizOpps jumps sharply past the $25,000 threshold of award amounts. The share of contracts that are publicly solicited in the government platform increases

[16]Because of the preliminary nature of these estimates, we concentrate on providing compelling visual evidence in the form of RDD plots, and refer to the magnitudes implied by them. Given the sample size and the magnitude of the effects, all of the preliminary estimates discussed in the following sections are statistically significant. We leave the formal estimation of coefficient magnitudes and appropriate standard errors for when we implement the full adjusted RDD procedure.

from roughly 16% at or slightly below $25,000, to 37% right past this threshold.

Figure 5: Publicizing requirement and use of FedBizOpps

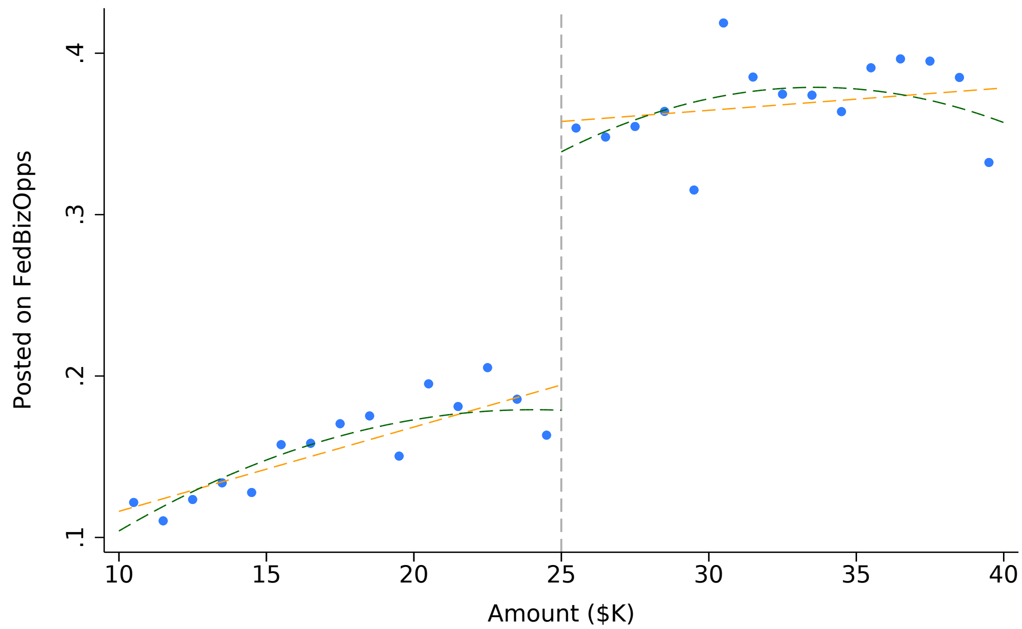

Notes: This figure shows the fraction of contracts posted on FedBizOpps by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Another way to see that the role of selection in this context is small is to consider the density of total awarded contracts again. As discussed previously, the total excess mass below the threshold due to contract manipulation is low, limiting the scope for bias. Additionally, it is also reassuring that pre-award contract characteristics do not vary sharply at the threshold, as we show in Appendix Figure A3.

Reassured by all this evidence, we proceed to estimate reduced form effects (Equation (8)) of publicizing contracts on three sets of outcomes: the intensity of competition, winning vendor characteristics (including its relationship with the awarding office), and post-award performance. Most of the existing literature has studied these variables independently.[17] By studying them jointly, we can generate a comprehensive understanding of the mechanisms and implications of policies oriented to enhance competition.

[17]See, for example, Athey (2001); Li and Zheng (2009) (competition), Macleod and Malcomson (1989); Bajari et al. (2009); Malcomson (2012) (relations), and Bajari et al. (2014); Decarolis et al. (2018); Ryan (2020) (ex-post renegotiation and performance).

4.2.1. Competition: Figure 6 shows how posting solicitations on FedBizOpps impacts the number of offers that a contract receives around the threshold. Contracts right above $25,000 (which are more likely to be publicly solicited), receive roughly 0.33 more bids. This effect is quite significant for this context: the average contract receives less than 2.5 offers below the threshold, indicating a reduced form effect of more than 13%. Moreover, since the policy only changes the likelihood of a publicized solicitation by around 24p.p. (see Figure 5), this implies a causal effect of approximately 67%.

Figure 6: Publicity and intensity of competition

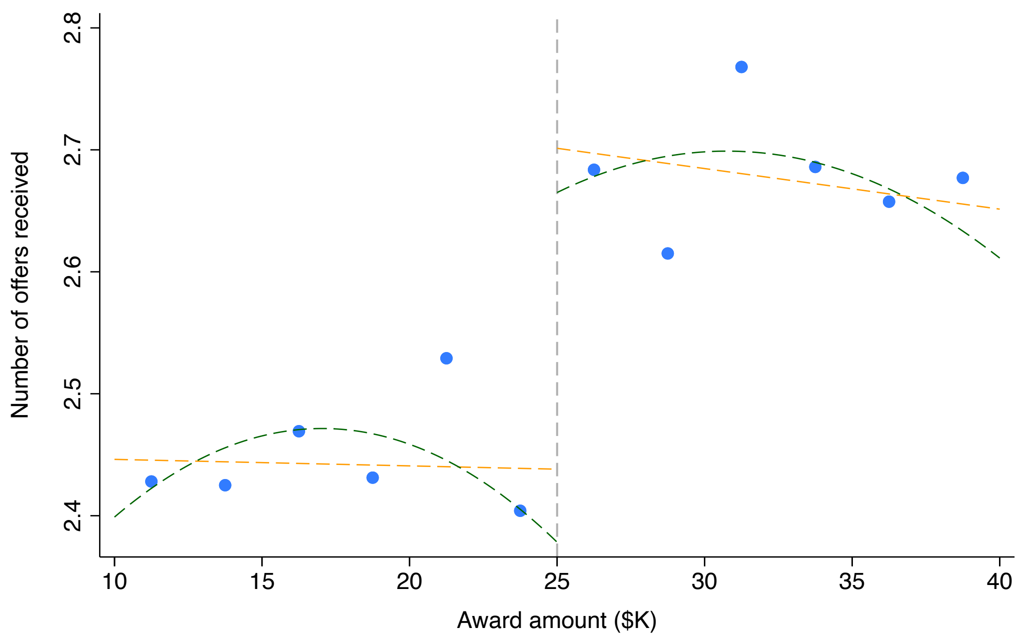

Notes: This figure shows the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data comes from FBO.gov and the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2019. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

These results indicate that encouraging the public posting of solicitations leads to the stated goal of increasing competition by attracting additional bids. However, an open question is whether these new offers affect the equilibrium allocation of the contract. We now turn to this question by analyzing the effects on the characteristics of the selected supplier.

4.2.2. Awarded Contractor Characteristics: Publicizing solicitations could attract marginal bidders that are not competitive, leaving the characteristics of the average selected supplier unchanged. However, Figure 7 shows that this is not the case. In Panel (a), we see that publicized contracts are awarded to vendors that have fewer prior contact with the office, as measured by the share of previous contract dollars for the same product category that was awarded to the firm. We also see in Panel (b) that publicized contracts are awarded to geographically more distant vendors relative to the location of the contracting office. Along the same lines, Panel (c) and Panel (d) show that these selected contractors are more likely to be foreign and less likely to be in the same state as the buyer.

Figure 7: Publicity and the characteristics of the winning firm

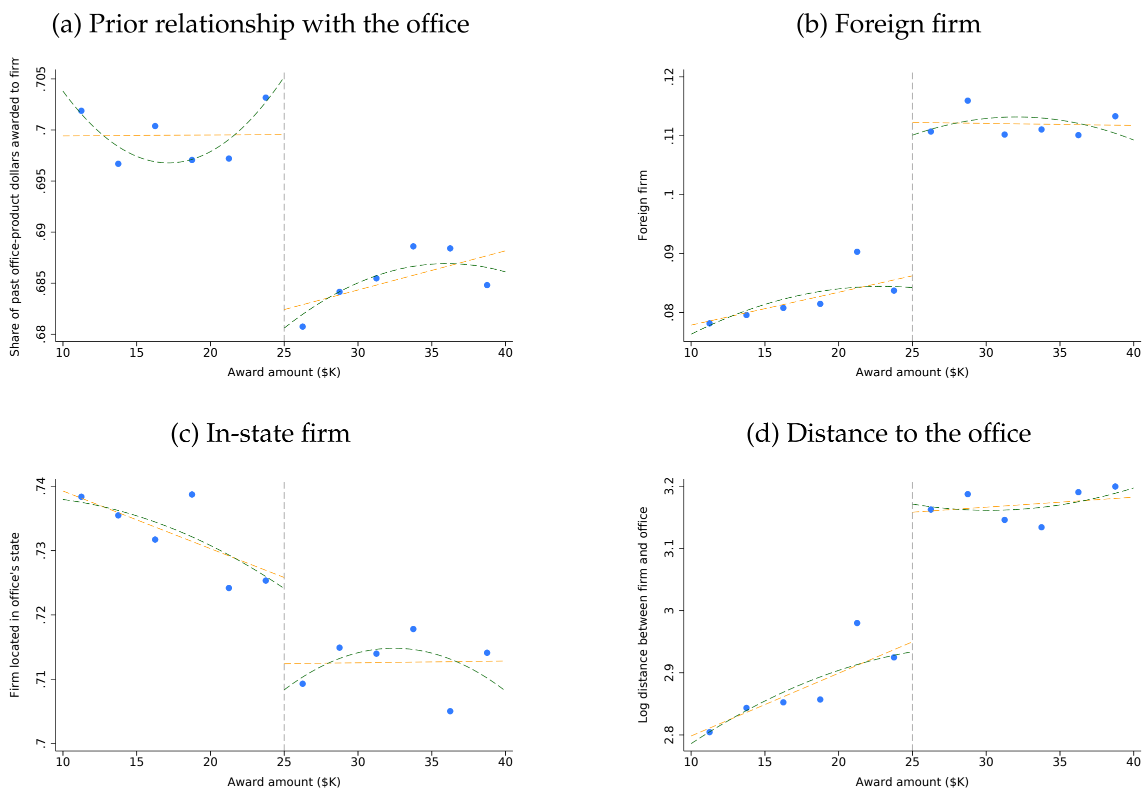

Notes: This figure presents four binned scatter plots, which depict an average outcome by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The outcome in each Panel is as follows: (a) the share of previous contract dollars by the same office and for the same product code, that had been awarded to the same winning firm; (b) an indicator equal to one if the contract is awarded to a foreign vendor; (c) an indicator equal to one if the contract is awarded to a firm in the same state as the contracting office; (d) the natural logarithm of the distance (in miles) from the contracting office's location and the vendor location. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2019. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Taken together, these results suggest that marginal entrants attracted by the public solicitation do win awards with a positive probability. Furthermore, compared to the vendors of unpublicized contracts, they tend to have less previous contact with the contracting office and are further away from their physical location.

4.2.3. Performance: We now measure effects on post-award performance. We use two measures of contract performance that are commonly used in the literature: cost overruns and delays (e.g. Kang and Miller, 2017; Decarolis et al., 2018; Carril, 2020). For each of our two variables, we consider two continuous measures—overrun dollars as a share of the original award and days of delay relative to expected schedule— and two dichotomic measures—any overruns or delays.

Figure 8 presents the results. We find an increase in all these four indicators of poor performance. The share of contracts with overruns and the share of contracts with delays both increase by 1p.p., which constitutes an effect of roughly 10%. Again, because the first-stage effect is around 15p.p., the causal IV estimate is in the order of 67%.

Figure 8: Publicity and post-award contract performance

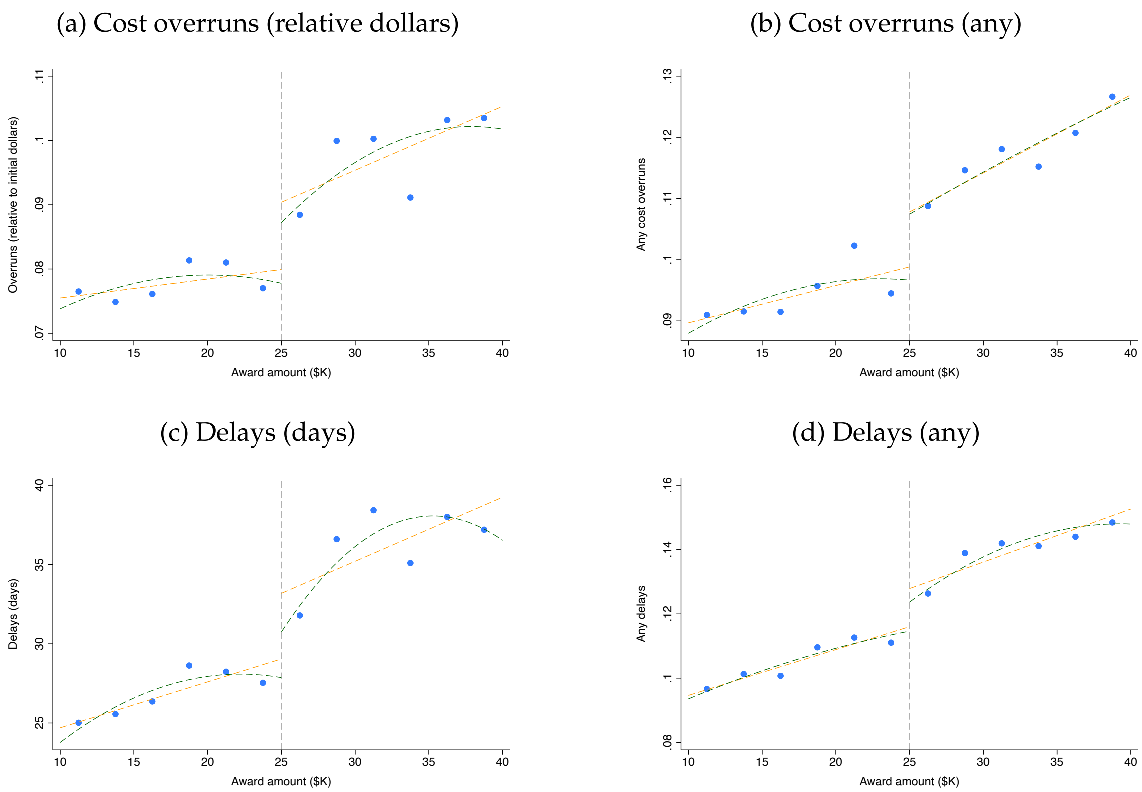

Notes: This figure presents four binned scatter plots, which depict an average outcome by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The outcome in each Panel is as follows: (a) the difference between actual obligated contract dollars and expected total obligations at the time of the award (i.e. cost overruns), divided by the expected dollar obligations at the time of the award; (b) an indicator equal to one if the contract has positive cost overruns; (c) the difference between actual days of contract duration and expected days of duration at the time of the award (i.e. delays); (d) an indicator equal to one if the contract has positive delays. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2019. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

4.2.4. Heterogeneity Analysis: We now examine whether the results presented above vary systematically across important contract characteristics. We start by considering differential effects for each of the three main Departments within the DOD. We do this by re-estimating our RDD specifications separately for the Army, the Navy, and the Air Force. Appendix Figure A4, Appendix Figure A5, and Appendix Figure A6 present the results for the number of offers, the probability of a foreign contractor, and cost overruns, respectively. While the increase in the number of offers seems to be present across all three agencies, the effects on foreign contractors and cost overruns seem to be mostly driven by the Army.

We then consider heterogeneous effects for goods and services. Appendix Figure A7, Appendix Figure A8, and Appendix Figure A9 respectively present the same three outcomes: number of offers, probability of a foreign contractor, and cost overruns. We see that publicizing contracts increases competition for both goods and services by a similar amount, despite the lower overall number of offers received in service contracts. A similar pattern emerges from the effects on the share of foreign firms: both goods and services see an increase due to the publicity requirements, even though their baseline levels differ. Finally, we see a different pattern for the effects on post-award performance. We see no change at the threshold for goods, and a significant deterioration of performance for services.

This last result is interesting because it is consistent with the theory of contractual incompleteness, the foundation of most concerns about the adverse effects of competition on post-award performance. Some transactions are easy to specify ex-ante, while others are more complex and involve countless possible contingencies. Writing down the precise terms for the purchase of a commodity (e.g., gasoline) is easier than to contract for an ad-hoc service, where post-award renegotiation and adaptation are likely to occur due to unspecified contingencies. Furthermore, bid evaluation for commodity contracts is likely a more straightforward endeavor than assessing the potential quality of bids for services. The results from Appendix Figure A9 are consistent with this intuition, finding that services and not goods drive negative performance effects.

To assess this mechanism more directly, we leverage the rich heterogeneity of our data, which features 1,918 distinct product categories. While some of them rarely experience performance issues ex-post, these are widespread for others. We use this fact to construct a proxy of complexity by product category. The proxy is based on the baseline level of post-award performance, which we define as the average performance experienced by contracts below $20,000. Based on this measure, we divide product categories into quartiles and estimate our RDDs separately for each of the four groups. If the contract incompleteness hypothesis is correct, then the negative impact of competition on post-award performance will be driven by those product categories that are more complex, defined as those that are exposed to negative performance even when not publicized.

Figure 9 supports this interpretation. We see that the average increase in overruns that we reported in Figure 8 is driven entirely by goods and services in the top quartile of complexity. We are unable to reject 0 for the lower three quartiles. These results are consistent with contract incompleteness being an important driver of our results.[18]

Figure 9: Heterogenous effects by contract complexity

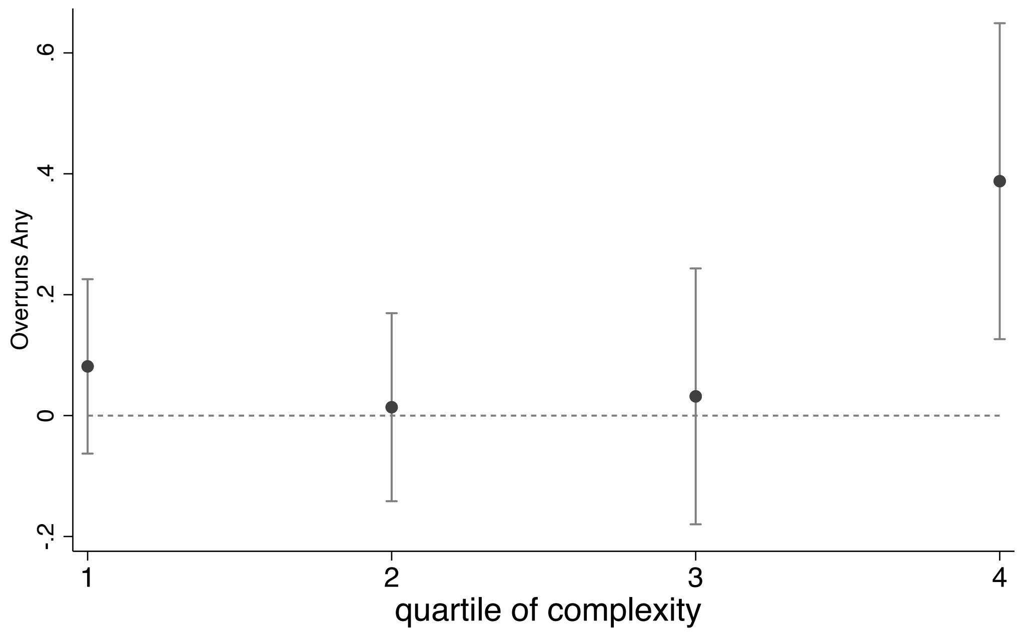

Notes: This figure shows four regression coefficients and their 95% confidence intervals. Each coefficient is a regression discontinuity estimate of the reduced-form Equation (8), estimated separately on four subsamples of the data. The dependent variable is an indicator for any positive cost overruns. The subsamples are determined by the four quartiles of a proxy of contract complexity. The contract complexity proxy is constructed at the product category level, and is defined as the average cost overruns for contracts with awards below $15,000 in that category. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017.

4.3. Comparing Pre-award and Post-award Effects

The results presented in the previous section imply that there is a trade-off associated with increasing competition for contracts by publicly disseminating information about them. On the one hand, competitive pressure leads to lower acquisition prices. On the other hand, selected suppliers are systematically different: they have less experience dealing with the contracting officer, and they perform worse ex-post.

How to weigh these different components ultimately depends on how buyers value each dimension. This motivates future work to try to recover buyers' preferences from their observed

[18] An issue with this test is that it uses an outcome measure to define the different groups of analysis. We only use contracts below $20,000 precisely to avoid any effect of the treatment on the classification, but the ideal measure would be based on pre-award characteristics of the contract. We are currently working on using data from the text of the solicitation to generate a better proxy of ex-ante complexity.

choices. By imposing additional structure to the key primitives of the setup in Section 3.1, we hope to use the revealed choices of publicizing decisions along with the exogenous variation generated by the regulation to estimate these parameters. Measuring how buyers assess the different procurement outcomes will speak to the severity of agency problems across the DOD contracting units. Furthermore, it can help us evaluate the consequences of counterfactual policies that affect competition and solicitation transparency.

In the absence of this analysis, it may still be informative to compare different dimensions of this trade-off. In particular, we can compare the ex-ante price reductions with the ex-post increases in cost overruns, since both of these are measured in the same units (dollars of spending). We can then ask if, on average, cost savings associated with more intense competition justify additional disbursements after the award. We make this comparison separately for 14 product categories—7 goods and 7 services—that have at least 2,000 contracts in our sample. For each category, we estimate average price effects (ex-ante) and average overruns effects (ex-post).

Figure 10 presents the results. Two key points are worth highlighting. First, the negative slope of the average effects is consistent with the trade-off we have emphasized so far: larger price reductions ex-ante are associated with larger increases in overruns. Second, in terms of magnitudes, the ex-ante price reductions are small relative to the ex-post cost increases. For all the categories analyzed, the dots in Figure 10 fall to the right of the 45 degree line, indicating that the net cost changes (cost overruns effects - price effects) are all positive.

Of course, these results do not provide a comprehensive assessment of the trade-off involved because buyers may value other dimensions not included in this analysis. For example, the publication of the offer on FedBizOpps may reduce the risk of a protested award by increasing transparency. However, the results constitute clear evidence that, at least in terms of resources spent in procurement, the policy seems to be backfiring: it generates modest cost savings at the time of the award, which are more than compensated by later cost increases.

From a policy perspective, given the volume of contracts impacted by this regulation, the consequences of the studied policies are meaningful. Specifically, the price range included our sample cover 65% of non-R&D, stand-alone contracts signed by the DOD. The price reductions due to enhanced competition generate annual cost savings in the order of $27.4 million. However, the increases in dollars spending due to increased overruns are in the order of $ 138.8 million per year. As a result, imposing publicity requirements aimed at enhanced competition and broaden industry participation result in excess in annual spending of $ 111.4 million.

Figure 10: Comparing ex-ante price effects with ex-post cost overruns

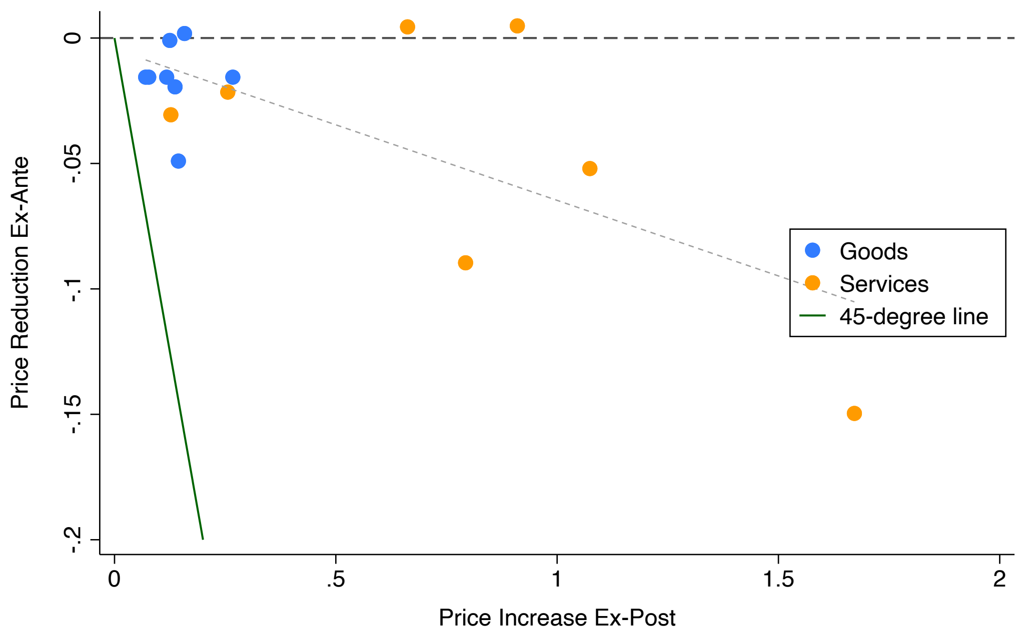

Notes: This figure shows the relationship between estimates of price effects (vertical axis) and cost overruns (horizontal axis) for 14 product categories that have at least 2,000 contracts in our sample. Price effects estimates are discussed in Section 4.1, while cost overruns estimates are discussed in Section 4.2. The green 45-degree line represents points where the (ex-ante) price reductions are equal to the (ex-post) cost overruns. We depict estimates for goods categories in blue and for service categories in orange. The data source is the Federal Procurement Data System-Next Generation. The sample consists of definitive contracts and purchase orders in selected categories with more than 2,000 observations, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017.

5. Conclusion

This paper studies the relationship between competition and public procurement outcomes. We do so by leveraging a regulation that generates quasi-experimental variation in the extent to which contract opportunities are broadly advertised to potential suppliers.

Our results emphasize that increasing competition via publicizing solicitations involves a trade-off for buyers. While the added competitive pressure results in lower acquisition prices, broader dissemination leads to a different pool of vendors, who on average perform worse ex-post. While a full evaluation of this trade-off requires more information about the government's preferences over different contract outcomes, we provide evidence that the publication requirement leads to increased procurement costs: price savings at the time of the award are smaller than observed increases in post-award cost overruns. We argue that the negative impact of competition on ex-post performance is explained by contract incompleteness. We do so by providing evidence that negative performance impacts are driven by contracts for goods and services that are relatively complex.

The estimated results open a number of questions that we pursue to answer in upcoming versions of the paper. In particular, we will propose a procurement model to investigate two interesting data patterns: First, the negative consequences of enhanced competition in this context is explained by the entry of a relatively small number of bidders that bid aggressively at the bidding stage, but perform poorly afterward, leading to frequent renegotiations and delays. We aim to test the implications of possible remedies studied by economic literature to address the adverse selection of bidders. Second, the buyer likely has private information about contractor's expected outcomes, we will propose an explicit choice framework to explain how buyers make publication decisions in this context, to recover their preferences and evaluate policy-relevant counterfactuals.

An important lesson of our analysis is that the effects of policies that seek to increase competition may have highly heterogeneous effects depending on the specific characteristics of the purchase. This is somewhat at odds with the largely homogeneous set of rules contained in the Federal Acquisition Regulation. Our results suggest that having policies that explicitly recognize and accommodate this heterogeneity could yield large increases in the efficiency of the public procurement system.

References

Athey, S. (2001). Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information. *Econometrica 69*(4), 861–889.

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

Bajari, P., S. Houghton, and S. Tadelis (2014). Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. *American Economic Review 104*(4), 1288–1319.

Bajari, P., R. McMillan, and S. Tadelis (2009). Auctions Versus Negotiations in Procurement: An Empirical Analysis. *Journal of Law, Economics, and Organization 25*(2), 372–399.

Bandiera, O., A. Prat, and T. Valletti (2009). Active and passive waste in government spending. *American Economic Review 99*(2009 Sept), 1278–1308.

Best, M. C., J. Hjort, and D. Szakonyi (2017). Individuals and Organizations as Sources of State Effectiveness. *NBER Working Paper No. 23350*, National Bureau of Economic Research.

Carril, R. (2020). Rules Versus Discretion in Public Procurement. *Mimeo*.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records. *Quarterly Journal of Economics 126*(2), 749–804.

Coviello, D. and S. Gagliarducci (2017). Tenure in Office and Public Procurement. *American Economic Journal: Economic Policy 9*(3), 59–105.

Coviello, D. and M. Mariniello (2014). Publicity Requirements in Public Procurement: Evidence from a Regression Discontinuity Design. *Journal of Public Economics 109*, 76–100.

Decarolis, F., L. M. Giuffrida, E. Iossa, V. Mollisi, and G. Spagnolo (2018). Bureaucratic Competence and Procurement Outcomes. *NBER Working Paper No. 24201*, National Bureau of Economic Research.

Gerard, F., M. Rokkanen, and C. Rothe (2016). Bounds on Treatment Effects in Regression Discontinuity Designs with a Manipulated Running Variable. *Nber Working Paper Series Bounds* (22892), 1–64.

Hart, O. and J. Moore (1988). Incomplete Contracts and Renegotiation. *Econometrica 56*(4), 755–785.

Kang, K. and R. A. Miller (2017). Winning by Default: Why is there so Little Competition in Government Procurement. *Mimeo*.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics 8*(1), 435–464.

Kleven, H. J. and M. Waseem (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan. *Quarterly Journal of Economics 128*(2), 669–723.

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

Laffont, J.-J. and J. Tirole (1990). Adverse Selection and Renegotiation in Procurement. *The Review of Economic Studies 57*(4), 597–625.

Li, T. and X. Zheng (2009). Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions. *Review of Economic Studies 76*(4), 1397–1429.

Li, T. and X. Zheng (2012). Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions. *Journal of Econometrics 168*(1), 29–46.

Liebman, J. B. and N. Mahoney (2017). Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement. *American Economic Review 107*(11), 3510–3549.

Macleod, W. B. and J. M. Malcomson (1989). Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. *Econometrica 57*(2), 447–480.

Malcomson, J. M. (2012). Relational incentive contracts. *The Handbook of Organizational Economics* (508), 1014–1065.

Ryan, N. (2020). Contract Enforcement and Productive Efficiency: Evidence From the Bidding and Renegotiation of Power Contracts in India. *Econometrica 88*(2), 383–424.

Saez, E. (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy 2*(3), 180–212.

Spulber, D. F. (1990, oct). Auctions and Contract Enforcement. *The Journal of Law, Economics, and Organization 6*(2), 325–344.

Williamson, O. E. (1976). Transaction-Cost Economics: The Governance of Contractual Relations. *Journal of Law & Economics 22*(2), 233–261.

Appendix

A. Additional Figures

Figure A1: Contract Timeline and Data Sources

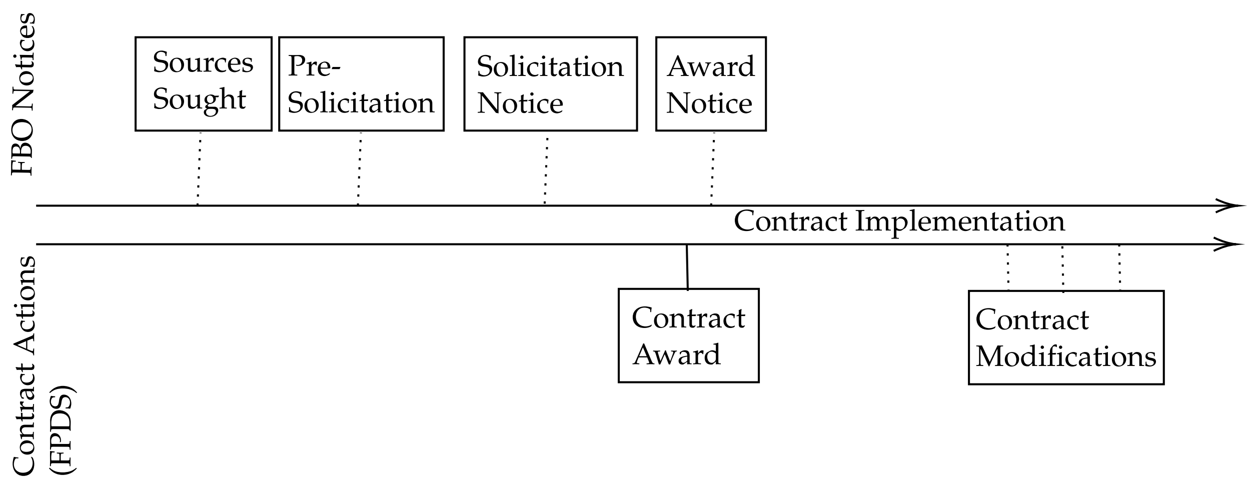

Notes: This figure presents a timeline of events associated with a typical contract. Milestones located above the arrows correspond to notices that are published on the government's point of entry (`fedbizopps.gov`). Milestones below the arrows generate information that is recorded on the Federal Procurement Data System (FPDS) - Next Generation.

Figure A2: Contract award density

(a) Raw density

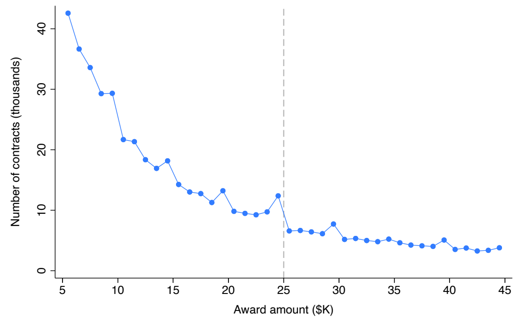

(b) Round number correction

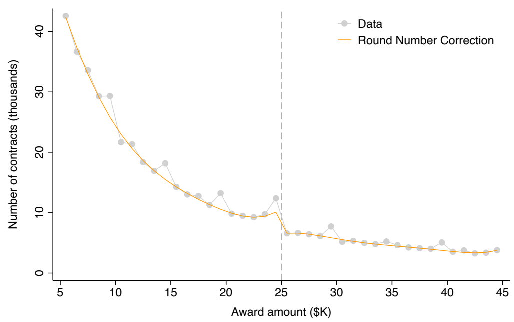

(c) Polynomial interpolation

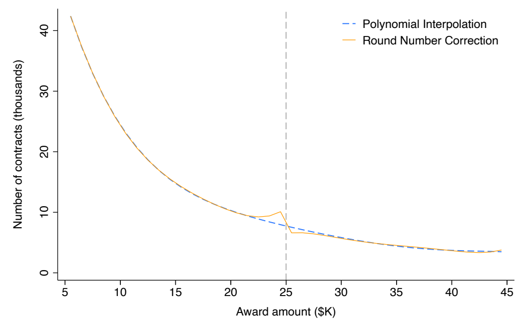

Notes: This figure shows the frequency distribution of contract awards as obtained from the data, as well as some non-parametric corrections. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Award amounts are discretized into right-inclusive bins of $1,000 dollars length. Panel (a) shows the raw frequency distribution. Panel (b) compares this with a smoothed function that corrects for round number effects at every $5,000 and $10,000 multiple, following the methodology from Kleven and Waseem (2013). Panel (c) compares the round number correction with a 5th-degree polynomial interpolation that excludes awards between $20,000 and $30,000.

Figure A3: Pre-award characteristics around the threshold

(a) Competitive award

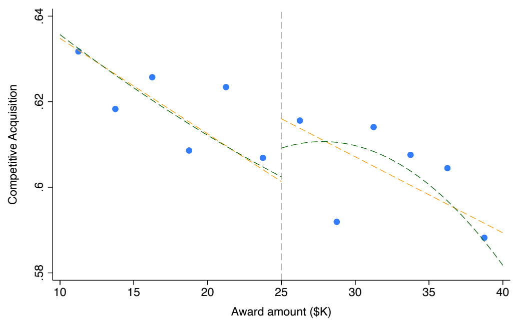

(b) Set-asides

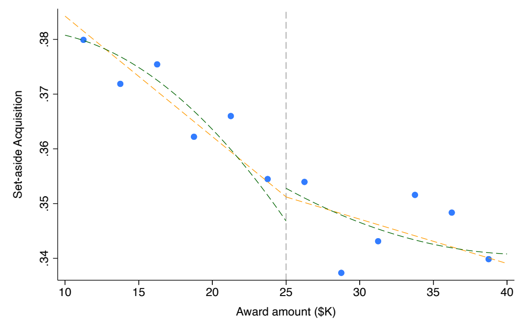

(c) Simplified acquisition procedures

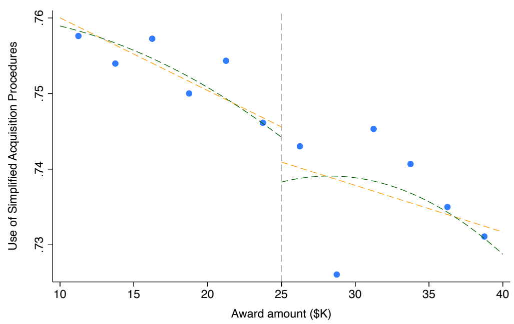

(d) Service contracts

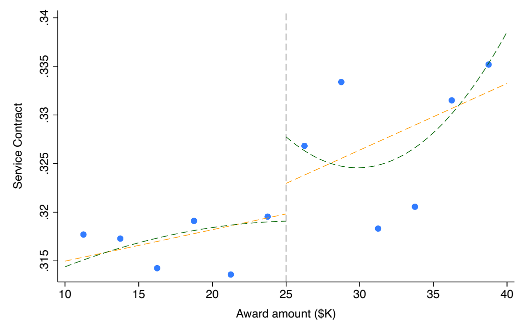

Notes: This figure presents four binned scatter plots, which depict an average pre-award characteristic by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The pre-award characteristic in each Panel is as follows: (a) an indicator equal to one if the contract was competitively solicited; (b) an indicator equal to one if the contract was set-aside for a preferential group (e.g. small businesses); (c) an indicator equal to one if the contract was awarded using simplified acquisition procedures; (d) an indicator equal to one if the award is for a service contract. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A4: Heterogeneous effects on competition by major departments

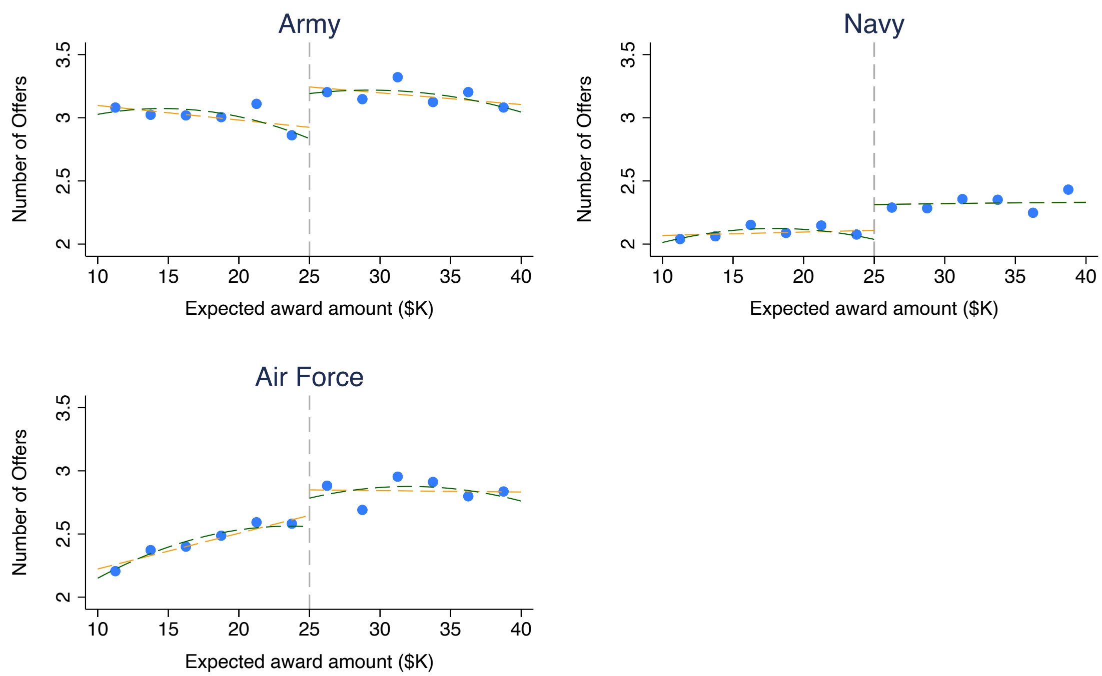

Notes: This figure presents three binned scatter plots, which depict the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards made by the Army. Panel (b) restricts the sample to awards made by the Navy. Panel (c) restricts the sample to awards made by the Air Force. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A5: Heterogeneous effects on winner characteristics by major departments

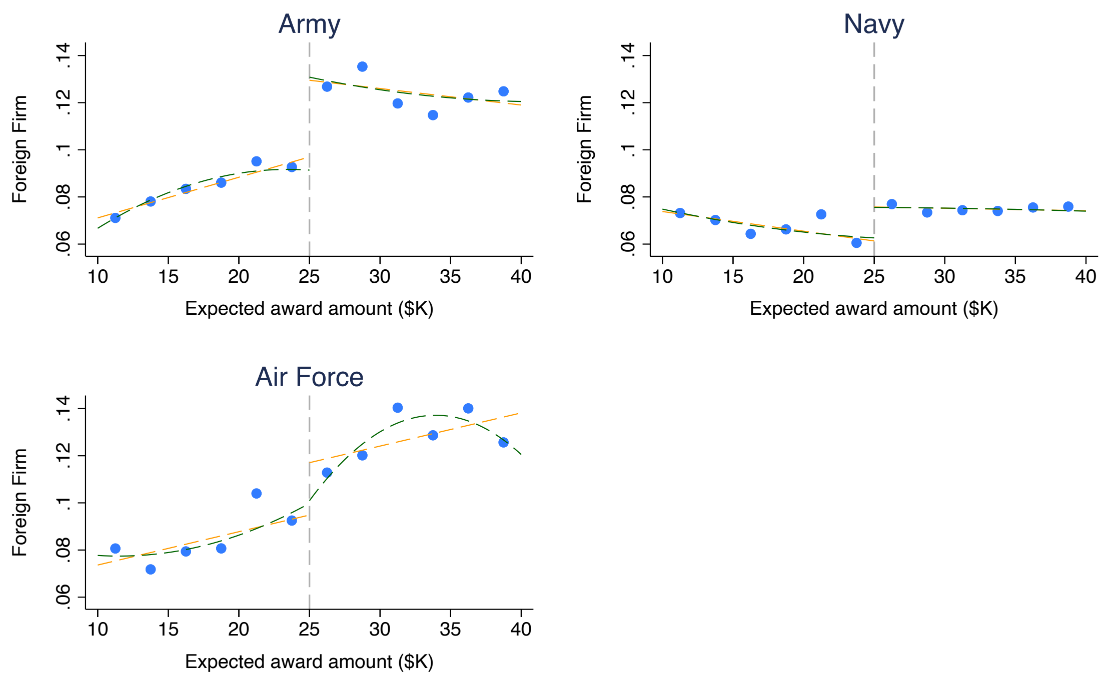

Notes: This figure presents three binned scatter plots, which depict the share of contracts awarded to a foreign firm by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards made by the Army. Panel (b) restricts the sample to awards made by the Navy. Panel (c) restricts the sample to awards made by the Air Force. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A6: Heterogeneous effects on performance by major departments

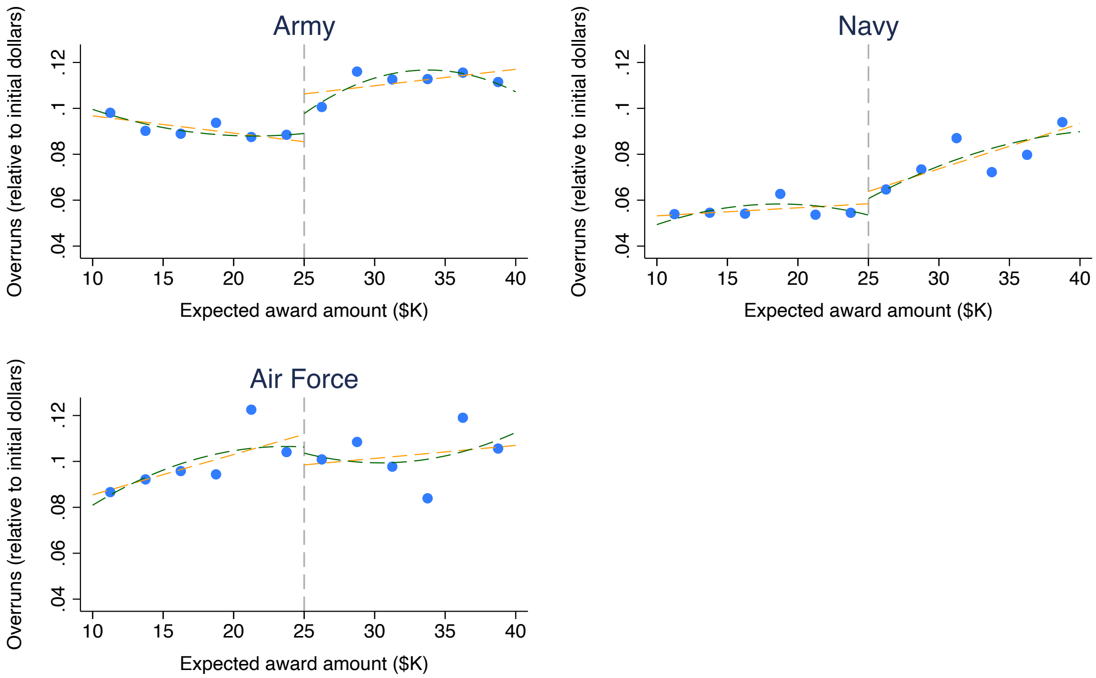

Notes: This figure presents three binned scatter plots, which depict average cost overruns by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the time of the award, divided by expected obligations. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards made by the Army. Panel (b) restricts the sample to awards made by the Navy. Panel (c) restricts the sample to awards made by the Air Force. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A7: Heterogeneous effects on competition: goods versus services

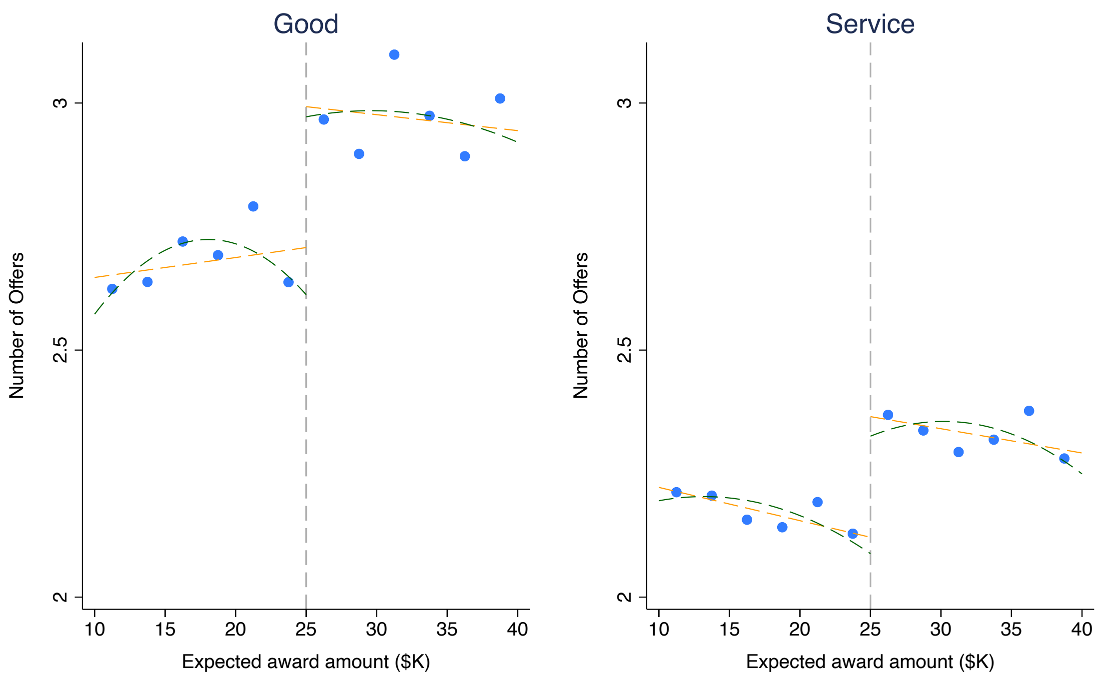

Notes: This figure presents two binned scatter plots, which depict the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A8: Heterogeneous effects on winner characteristics: goods versus services

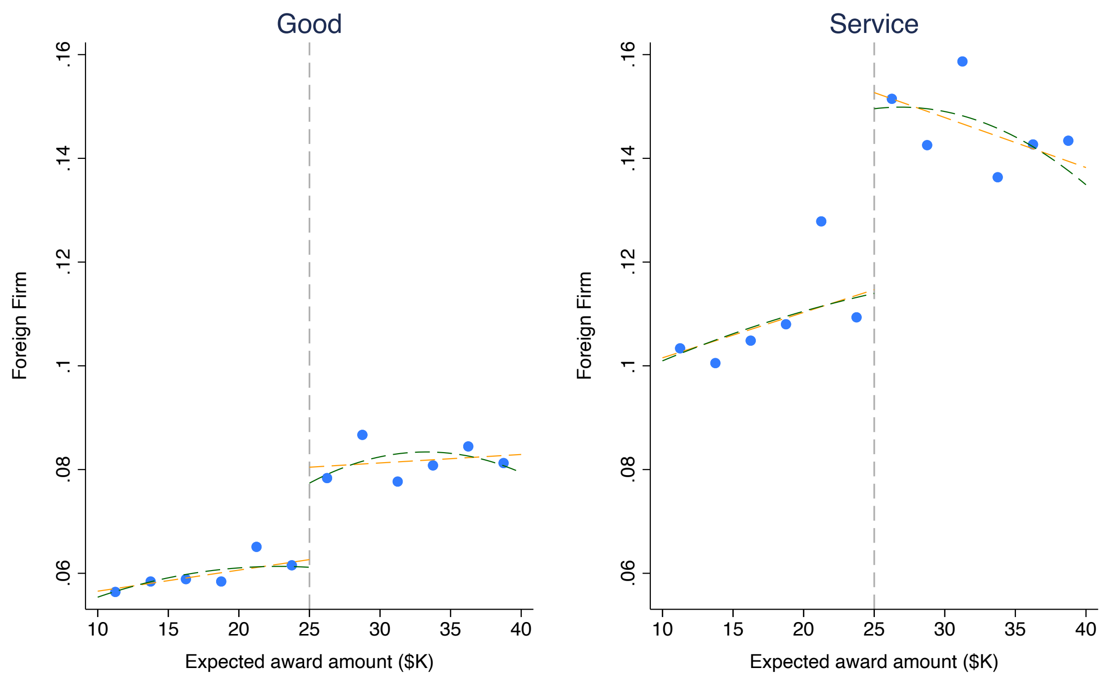

Notes: This figure presents two binned scatter plots, which depict the share of contracts awarded to a foreign firm by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A9: Heterogeneous effects on performance: goods versus services

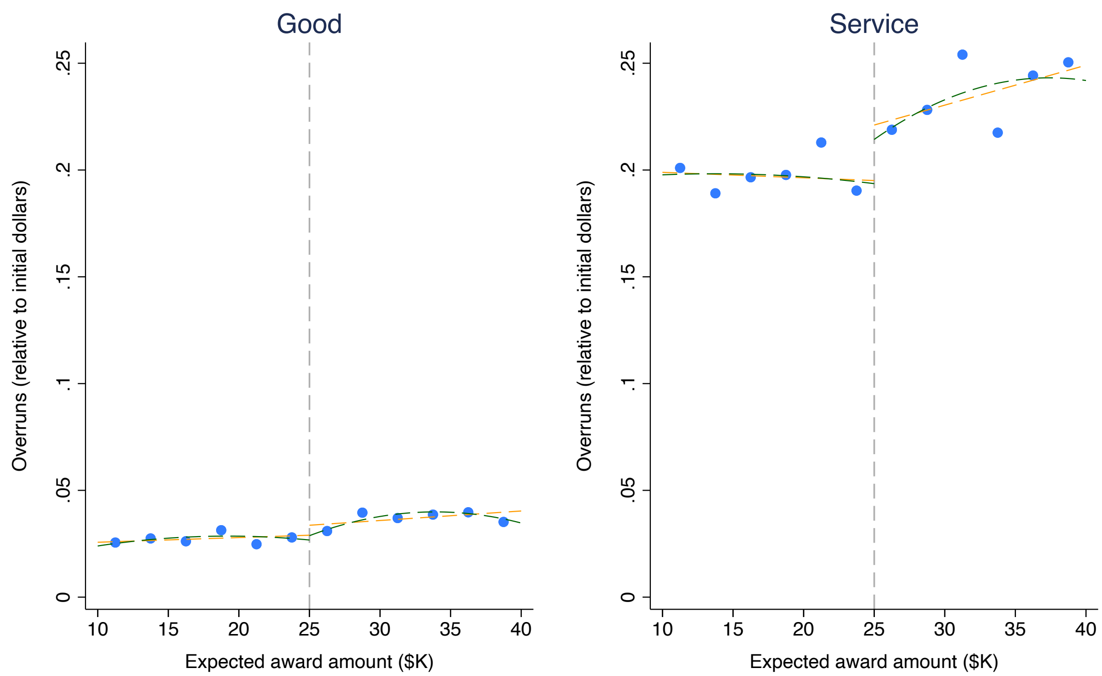

Notes: This figure presents two binned scatter plots, which depict average cost overruns by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the time of the award, divided by expected obligations. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 5,000 and $ 45,000, awarded by the Department of Defense in fiscal years 2011 through 2017. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A10: Model fit

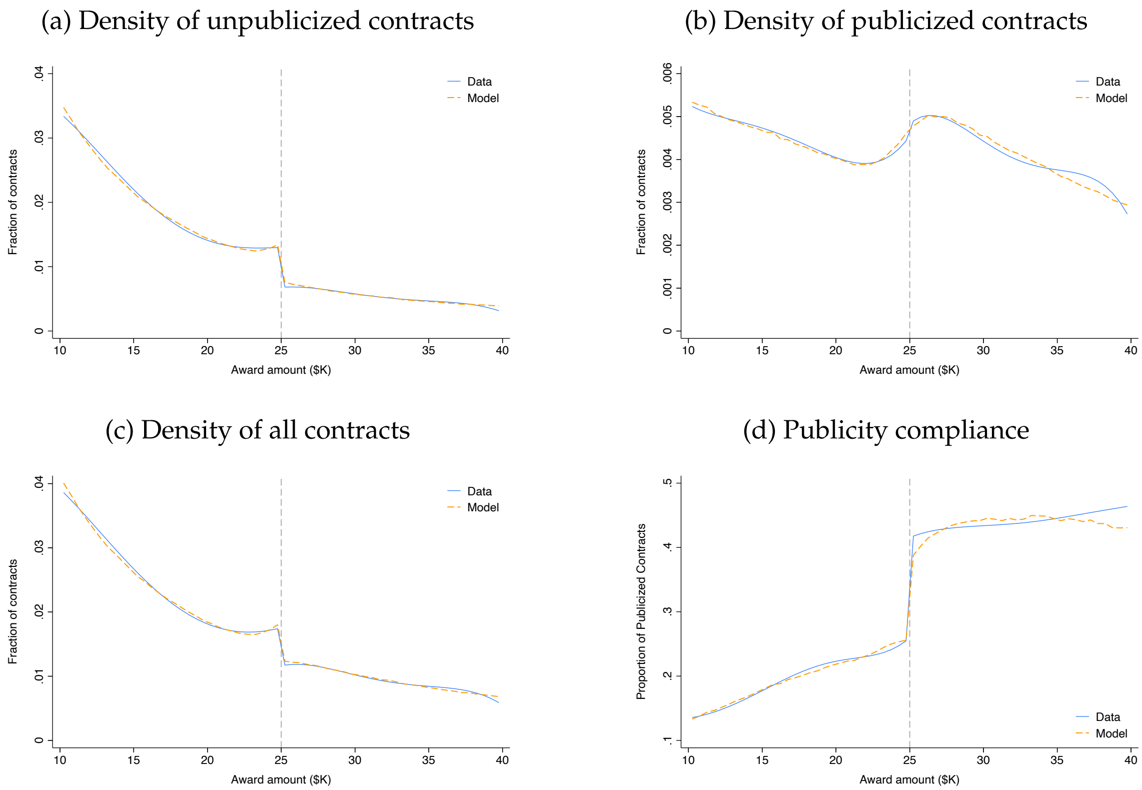

Notes: This figure presents the model fit, based on a simulated method of moments estimation. In each panel, moments based on (smoothed) data are presented in solid blue lines, while model-based simulated moments are presented in dashed orange lines. Panel (a) presents the density of unpublicized contracts, Panel (b) the density of publicized contracts, Panel (c) the total density (the sum of (a) and (b)), and Panel (d) presents the share of publicized contracts at each award level.

B. Additional Tables