American Economic Review 2026, 116(2): 535–581
https://doi.org/10.1257/aer.20221345

Competition under Incomplete Contracts and the Design of Procurement Policies†

By Rodrigo Carril, Andres Gonzalez-Lira, and Michael S. Walker*

We study the effects of intensifying competition for contracts in the context of US Defense procurement. Leveraging a discontinuous regulation that mandates agencies to publicize certain contract opportunities, we document that expanding the set of bidders reduces award prices but deteriorates post-award performance in terms of cost overruns and delays. We develop and estimate an auction model with endogenous entry and stochastic execution performance, in which the buyer endogenously chooses the intensity of competition. Model estimates indicate substantial heterogeneity in performance across contractors and show that simple adjustments to the current regulation could provide significant savings in procurement spending. (*JEL* D22, D24, D44, D82, D86, H56, H57)

Buyer-seller transactions—concerning everything from standardized goods, such as office supplies or fuels, to customized needs, such as construction projects or consultancy services—are often governed by competitively awarded procurement contracts. The pervasive use of competition to assign contracts stems from the notion that competitive bidding can be a powerful tool to reduce procurement prices (Bulow and Klemperer 1996). Yet, more intense competition for contracts that involve customized obligations and deliverables could allow underqualified contractors to win, leading to deficient execution ex post. Therefore, the assessment

*Carril: Universitat Pompeu Fabra and Barcelona School of Economics (email: rodrigo.carril@upf.edu); Gonzalez-Lira: Pontificia Universidad Católica de Chile, Business School (email: andresgonzalezlira@uc.cl); Walker: US Department of Defense (DOD) (email: michael.s.walker110.civ@mail.mil). Sylvain Chassang was the coeditor for this article. This paper was the main chapter of Gonzalez-Lira's PhD dissertation. He thanks his advisors Steve Tadelis, Benjamin Handel, Kei Kawai, and Reed Walker for their guidance and support. We thank our discussants, Francesco Decarolis, Paulo Somaini, Karam Kang, Decio Coviello, and Alex Arsenault-Morin for many helpful comments. We also thank Christopher Campos, Ernesto Dal Bo, Mark Duggan, Liran Einav, Ying Fan, Matt Gentzkow, Phil Haile, Jon Kolstad, Fernando Luco, Cristobal Otero, Giancarlo Spagnolo, Dario Tortarolo, Damian Vergara, Christopher Walters, Heidi Williams, Guo Xu, Roman David Zarate, Ph.D. students of the UC Berkeley research methods course, conference and seminar participants at Stanford, UC Berkeley, Yale, and several other institutions for valuable comments and suggestions. Gonzalez-Lira gratefully acknowledges financial support from the Institute for Business Innovation at the Haas School of Business. Carril gratefully acknowledges financial support from the Bradley Graduate and Postgraduate Fellowship through a grant to the Stanford Institute for Economic Policy Research, the Institute for Research in the Social Sciences at Stanford University, and the Spanish Agencia Estatal de Investigación through the Juan de la Cierva award (FJC2021-047328-I AEI/MCIN/EU/PRTR), research grant PID2020-115044GB-I00/MICIU/AEI/10.13039/501100011033, and the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). The views expressed herein are those of the authors and do not reflect the position of the Department of the Army or the DOD.

†Go to https://doi.org/10.1257/aer.20221345 to visit the article page for additional materials and author disclosure statement(s).

of competition promotion should account for both potential benefits due to contract price reductions and possible adverse effects due to poor execution.

An empirical investigation of this trade-off is complicated due to the need for comprehensive data on contract execution and a compelling research design. In this paper, we make progress on both fronts to study the equilibrium effects of enhancing competition for procurement contracts in acquisition prices and execution performance. We focus on DOD procurement, a setting of relevance given that it awards \$500 billion in contracts per year, representing more than half of federal procurement spending. Moreover, this setting provides us with policy variation in the degree of contract competition, as well as detailed administrative data throughout the life cycle of each DOD contract, from design through execution.

Our empirical strategy exploits regulation that requires agencies to publicize contract opportunities that are expected to exceed \$25,000 in value through a centralized online platform. Publicity increases competition by expanding the set of potential bidders informed about the auction. Analyzing contract awards between \$10,000 and \$40,000, we exploit the discontinuous nature of the publicity requirements to estimate its effects on four sets of outcomes: (i) the actual number of bids received, (ii) characteristics of the buyer-contractor relationship, (iii) procurement costs, and (iv) post-award contractor performance. By providing evidence on all of these fronts, we comprehensively characterize the consequences of changing the degree of competition for procurement contracts through this advertising channel. Furthermore, we exploit rich heterogeneity in the types of contracts that the DOD awards to assess the role of contract incompleteness in explaining our results.[1]

To estimate the price effects of contract publicity, we propose a method that recovers these effects from discontinuities in the conditional densities of publicized and nonpublicized contracts. We then estimate the effects of publicizing contract opportunities on three sets of nonprice outcomes—the number of bids received, the characteristics of the selected vendors, and post-award performance—using a regression discontinuity design (RDD). We find that contract awards advertised through the government platform see an increase in the number of bids of roughly 60 percent, confirming that the policy translates into a substantial increase in participation. We show that these marginal participants are competitive, leading to changes in the characteristics of winning firms: Awardees of publicized solicitations are, on average, 14 percent less likely to be small businesses and are located 60 percent farther from the buying agency. Furthermore, increased competition leads to contract price reductions: Publicized contracts are, on average, awarded at 6 percent lower prices. However, advertised contracts result in worse ex post performance: The probability of experiencing cost overruns and delays in the implementation stage increases by 7 percent and 8 percent, respectively. Performance results are driven mainly by service contracts—as opposed to goods purchases—and by contracts that

[1] Focusing on a window around the policy threshold implies that our sample does not include contracts related to major DOD acquisitions (e.g., fighter jets or weapon systems). However, the median contract awarded by the DOD is worth \$19,800, and 68 percent of contracts obligate less than \$40,000. Focusing on this range provides our results with a higher level of external validity, as contracts in our sample are much more similar to nondefense agencies and the private sector, compared to large-scale defense contracts. It is also worth noting that the volume of contracts impacted by the publicity regulation makes its implications economically meaningful. In 2018 alone, the DOD publicized contract solicitations valued at \$5.56 billion via the online platform FedBizOpps.gov.

we ex ante characterize as more complex. Finally, we find that these patterns are primarily accounted for by variation across firms (consistent with adverse selection) and not by within-firm responses to changes in the competitive environment.

Our reduced-form results suggest that promoting competition has mixed effects on contract outcomes: While it reduces the winning bid, it leads to worse outcomes at the execution stage. Suppliers' identities matter for explaining this variation in contract outcomes. Overall, promoting competition hinders buyers' ability to restrict participation to qualified vendors while attracting new participants who tend to perform poorly ex post. These results are stable across different estimation approaches and robustness checks that account for possible sources of bias.[2]

Motivated by this evidence, we develop and estimate an equilibrium model of competition for procurement contracts. The model allows us to estimate the underlying firm characteristics that shape adverse selection in this market and the role of buyer preferences in the decision to promote competition. Furthermore, we use the model to evaluate the consequences of counterfactual policies aimed at reducing public spending.

Our model consists of four stages that cover the different phases of a procurement project. First, a buyer decides on the degree of competition by choosing whether to openly publicize the contract or to invite only specific contractors. Second, firms that receive information about the contract simultaneously decide whether to prepare a bid. They do this by comparing the expected utility of participating with the idiosyncratic cost of preparing the bid. Third, each bidder submits a bid that depends on the realization of a production cost estimate, consisting of a private component and a common component, which accounts for unobserved heterogeneity (Krasnokutskaya 2011; Haile and Kitamura 2019).[3] The award mechanism is a first-price, sealed-bid auction. Fourth, the awarded contractor executes the contract. Execution performance depends on the existence and magnitude of cost overruns stemming from an idiosyncratic shock realized ex post. The model allows for a correlation between pre-award production cost realizations and post-award performance shocks and incorporates the potential asymmetry between bidders who are informed directly by the buyer and those who participate only when the contract is openly publicized. Moreover, the model does not impose restrictions on buyers' preferences over outcomes and allows for idiosyncratic preferences for certain vendors that are uncorrelated with contract outcomes.

We estimate our model using data on publicity choices, auction participation, contract prices, and observed cost overruns. We exploit variation in market structure and the publicity threshold to identify the model's parameters. Our estimates highlight an asymmetry between the sellers whom the buyer would invite directly in the absence of publicity and those who bid only when the solicitation is openly publicized. The added bidders have slightly lower production costs and substantially lower participation costs, which makes them more likely to participate ceteris paribus. They are also considerably more prone to experience cost overruns in the execution stage.

[2] Our RDD setting is characterized by measurement error in the assignment variable. We address this concern directly, yet ultimately find that it has a quantitatively modest effect on the results.

[3] Controlling for unobserved heterogeneity is important since, in the procurement setting, bidders likely have more information about the auctioned contracts than the econometrician.

For both types of bidders, we estimate a negative correlation between production cost realizations and cost overrun shocks, implying that winning more competitive auctions will be associated with lower bids and higher overruns. Still, most of the negative effects of publicity on performance are quantitatively accounted for by the differential selection of bidders, not by within-type changes in the distribution of overrun realizations. Finally, we estimate that buyers have a preference for lower prices, lower cost overruns, and incumbent suppliers.

By specifying the selection process that leads to buyers' publicity choices, the model allows us to extrapolate the local effects estimated at the policy discontinuity to the full sample. We then use our model parameters to estimate the impact of promoting competition through publicity under the current regulation, as well as under alternative policy scenarios. Overall, our findings are consistent with the estimated reduced-form effects. Increasing competition has heterogeneous effects, leading to cost reductions when the transaction unit is relatively simple. However, competition backfires when the contract involves a complex product category, as increases in cost overruns exceed price reductions at the award stage.

Our results show that imposing regulations to promote bidder participation involves a risk of allowing underqualified firms to bid. An alternative policy design may rely on buyers (i.e., each local agency) to decide whether to publicize each contract. Delegating this decision to the buyer involves a trade-off. On the one hand, more discretion allows the buyer to tailor decisions, mitigating the potential risks of intensifying competition. On the other hand, the buyer could use this added discretion opportunistically, restricting competition to favor specific contractors. We use our model to simulate equilibrium outcomes under a deregulated setting where the buyer decides whether to publicize each contract. We find that the effects of discretion are ambiguous and critically depend on the degree of complexity of the purchase. Relative to no-publicity or full-publicity rule benchmarks, discretion yields lower contract costs for an intermediate range of complexity values, suggesting ample space to improve the current regulation.

We use our model to identify improvements to the current policy design, which depart from uniform publicity requirements. Many policies that regulate competition in public procurement settings—including the one we study—are strikingly simple: They do not differ depending on whether the transaction involves a commodity or a highly customized service. This mismatch between unsophisticated policies and a highly diverse set of transactions suggests meaningful room for improvement in policy design. We study the effects of introducing publicity requirements tailored to the complexity of the purchase, thus leveraging the benefits of intense competition for simple products while limiting its adverse impact on complex products. We find that the cost-minimizing level of publicity for fully specified products is close to 100 percent, whereas more complex product categories should require low use of publicity. This reduces average defense procurement costs by 1.5 percent, or \$104 million annually.

This paper contributes to the literature studying the classic problem of procurement in the presence of noncontractible quality concerns. Many papers have theoretically emphasized the shortcomings of standard (price-only) auctions and have studied the properties of alternative mechanisms (e.g., Spulber 1990; Bajari 2001; Calzolari and Spagnolo 2009; Decarolis 2018), including the derivation of optimal mechanisms in models with adverse selection (Manelli and Vincent 1995; Lopomo, Persico, and Villa 2023) and moral hazard (Burguet, Ganuza, and Hauk 2012; Chillemi and Mezzetti 2014).[4] This paper relates to the empirical literature showing that awarding mechanisms aimed at promoting price competition can be outperformed by less competitive alternatives (e.g., Bajari, McMillan, and Tadelis 2009; Decarolis 2014; Conley and Decarolis 2016; Coviello, Guglielmo, and Spagnolo 2018). Unlike most previous work, our setting offers a source of exogenous variation in the intensity of competition while keeping the awarding mechanism and contract design fixed. This allows us to isolate the role of competition intensity on contract outcomes. Moreover, while existing papers typically focus on a single type of acquisition, our rich sample containing a wide range of product categories allows us to make a broader empirical contribution and show how the implications of promoting competition depend on the relative importance of adaptation costs.

Our reduced-form analysis builds on closely related work by Coviello and Mariniello (2014), who use an analogous threshold-based publicity requirement policy in Italy to study its effects on similar outcomes as the ones we consider. They find that publicity leads to an increase in the number of bids, a significant reduction in prices, and no effect on post-award performance as measured by delays. While there are many similarities between our settings—most significantly, the policy variation and the fact that auctions are sealed bid and price only—two important differences may limit the direct comparability of both sets of results. First, the mechanism used to award contracts in their setting is not a first-price auction but the so-called average bid auction (ABA).[5] Since the ABA does not select the lowest bidder but a bid that tends to be close to the average bid, we need not expect the same type of competitive effects of publicity. The authors acknowledge that the effect of increased participation in equilibrium ABA bids is theoretically ambiguous, even absent endogenous entry considerations.[6] Furthermore, by eliminating very low bids, the effect of competition on post-award performance also becomes more subtle.[7] A second relevant difference is that the size of the "treatment" in both settings differs substantially. Empirically, while publicity requirements in their setting increase the number of bidders by 3.3 from a mean of 36, we find an increase of 2.3 bids from a mean of 3.5.[8] This suggests that the shock to competition may be more pronounced in our context (Bresnahan and Reiss 1991).

We also contribute to a growing literature that evaluates policies aimed at promoting (or restricting) bidders' participation in procurement auctions. This literature emphasizes that expanding the pool of potential bidders may not necessarily

[4] Asker and Cantillon (2008, 2010) study cases where quality is contractible and can, therefore, be explicitly incorporated into the awarding rule. Dini, Pacini, and Valletti (2006) provides an overview of these theoretical arguments and discusses the issues related to the practical implementation of scoring rules. In our setting, even if some nonprice components were contractible, buyers are constrained by regulation to use price-only auctions.

[5] The ABA awards the contract to the lowest bidder only after excluding bids that are deemed to be "too low" (Albano, Bianchi, and Spagnolo 2006). Decarolis (2018, p. 396) describes the selection rule as follows: "Disregard the top and bottom 10 percent of the bids; calculate the average of the remaining bids (call it A1); then calculate the average of all the bids strictly above the disregarded bottom 10 percent and strictly below A1 (call this average A2); the first price above A2 wins the contract and is paid his own price bid to complete the work."

[6] Decarolis (2018) shows that this type of mechanism features multiple symmetric Bayesian-Nash equilibria and that equilibrium bids may increase and converge to the reserve price as the number of bidders grows large.

[7] The authors can replicate their key findings in a small sample of contracts awarded through first-price auctions, yet all the baseline analysis is conducted on contracts awarded through ABA.

[8] This is the instrumental variable's estimate (see Table 2 column 5). The reduced form effect is 0.36 (Table 1 column 1).

translate into lower award prices if bidders' participation is endogenous, as their equilibrium bidding behavior may become less aggressive (Athey et al. 2011, 2013; Li and Zheng 2009, 2012; Krasnokutskaya and Seim 2011; Marmer et al. 2013; Bhattacharya et al. 2014; Sweeting and Bhattacharya 2015).[9] We explore this issue by explicitly modeling entry and bidding decisions and leveraging variation in the number of potential bidders that stems from exogenous changes in publicity requirements. We indeed find that incumbents are less likely to participate when anticipating fiercer competition. However, the competitive effect from new entrants dominates that of less aggressive bidding by incumbents, reducing the winning bid as a result.

Finally, this paper belongs to the literature that examines buyers' discretion on public procurement outcomes. This line of work has focused on the incentives that buyers face and highlights that their actions can be motivated by objectives other than simple contract cost reductions (Bandiera, Prat, and Valletti 2009; Liebman and Mahoney 2017; Coviello and Gagliarducci 2017; Coviello, Guglielmo, and Spagnolo 2018; Best, Hjort, and Szakonyi 2023; Decarolis et al. 2025; Carril 2022; Szucs 2024). Our contribution to this literature is to emphasize how buyers' actions impact sellers' endogenous decisions to participate in procurement markets. In this respect, our work is related to Kang and Miller (2021), who use IT procurement contracts in the United States to estimate a principal-agent model in which the buyer exerts costly effort to determine the degree of competition that contracts receive. Our paper is also concerned with modeling how the level of competition is endogenously determined, but rather than emphasizing a trade-off between the price benefits of competition and the costly effort of promoting it, we highlight the tension between price competition ex ante and deteriorated performance ex post.

The rest of the paper is organized as follows. Section I provides background on the DOD procurement system and the data we use for our analysis. In Section II, we provide evidence on the effects of contract publicity on a range of relevant outcomes. In Section III, we develop and estimate an equilibrium model of procurement competition, which we then use to study outcomes under counterfactual environments in Section IV. Section V concludes.

I. Setting and Data

A. *US Federal Procurement and Publicizing Requirements*

Public procurement is a large component of the US economy. In fiscal year 2019, federal contract awards totaled $926 billion. Contracts are awarded at highly decentralized levels, with over 3,000 different contracting offices that are part of an executive or independent agency.[10] The workforce in charge of public contracting

[9] These ideas were initially introduced by Samuelson (1985) and Levin and Smith (1994). Li and Zheng (2009) provide an empirical framework highlighting that increasing the number of potential bidders within the independent private values setting has ambiguous effects since two counteracting effects occur in equilibrium: a "competition effect" and an "entry effect." The former tends to reduce prices, while the latter tends to increase them.

[10] Executive agencies are headed by a cabinet secretary, like the DOD, 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.

comprises over 35,000 contracting officers whose primary role is to plan, carry out, and follow up on purchases made by their units. The Federal Acquisition Regulation (FAR)[11] defines and limits the contracting officers' scope of action (48 CFR, 2025). The FAR lays out policy goals, guiding principles, and a uniform set of detailed policies and procedures to guide the procurement process. Our analysis leverages a specific section of the FAR that regulates how buyers should publicize contract actions as a convenient source of quasiexperimental variation.

FAR part 5 requires publicizing contract opportunities to "increase competition," "broaden industry participation," and "assist small businesses [and other minority businesses] (…) in obtaining contracts" (48 CFR pt. 5.002, 2025). Since October 1, 2001, contract actions that exceed $25,000 must be publicized on an online government-wide platform (48 CFR pt. 5.101, 2025), which we will refer to as FedBizOpps (FBO).

Officers with contracts not expected to exceed this threshold are not required to publicize in FBO; however, they are still free to use it to increase contract visibility.[12] The regulation also allows for exemptions to the requirement above the threshold if doing so "compromises national security," if "the nature of the file does not make it cost-effective or practicable," or if "it is not in the government's interest." Therefore, while this policy discretely affects the likelihood of publicizing contracts around the threshold, compliance may be far from perfect given the voluntary nature of the rule below this value and the availability of exceptions above. Importantly, no other relevant regulations vary at this same threshold value. Supplemental Appendix C.1 presents additional policy and website details.

The overwhelming majority of the contracts analyzed in this paper are allocated through so-called simplified acquisition procedures, with vendors being selected according to the lowest price quote that is technically acceptable given the specifications.[13] This simple lowest-bid mechanism contrasts with more sophisticated awarding procedures available for larger acquisitions, which may consider attributes of the offer other than price. For example, the collection and use of past performance information from the Contractor Performance Assessment Reporting System (CPARS) are mandated for contracts with award values that largely exceed those considered in our analysis.[14]

[11] *Code of Federal Regulations*, Title 48: "Federal Acquisition Regulations System." (2025).

[12] Procurement officers with contracts with expected values below the threshold are only required to advertise the solicitation "by displaying [it] in a public place" (48 CFR pt. 5.101-a2, 2025). This includes, for example, a physical bulletin board located at the contracting office.

[13] We conducted several interviews with contracting officers who confirmed that contract awards in this dollar range are virtually always awarded to the lowest quote and that there is little discretion to deem a particular offer unacceptable. The most common reasons are a blatant omission of the solicitation specifications or the fact that the contractor is debarred from conducting business with the government based on past experience.

[14] For the DOD as of May 2021, the thresholds are $5 million for systems and operations support; $1 million for services and information technology; $750,000 for construction; and $500,000 for ship repair and overhaul. The only exception is architect-engineer contracts, with a threshold of $35,000, but these acquisitions represent less than 0.1 percent of our main analysis sample.

B. *Data*

We use two complementary sources of data. The first consists of the historical files from FBO, which provide detailed information on pre-award notices (i.e., solicitations) posted on the platform (US General Services Administration 2007–2020). The second is the Federal Procurement Data System-Next Generation (FPDS-NG), which tracks federal contracts from the time of their award and includes all follow-on actions, such as modifications, terminations, renewals, or exercises of options (US Bureau of the Fiscal Service 2006–2019).

We merge awards from FPDS-NG to notices on FBO using the solicitation number. However, while FPDS-NG contains the universe of federal awards, FBO only has the notices posted on the website. From this matching process, we construct a dummy variable equal to 1 if we can merge a contract with any pre-award notice on FBO, in which case we say the contract was publicized. Supplemental Appendix C.2 provides additional details on the construction of the dataset. Supplemental 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, subagency, 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. Furthermore, we observe the reason for and content of all contract modifications after the contract is awarded. These actions often involve extending the duration and/or increasing the dollar amount allocated to the vendor. We further complement this dataset with information about contracting offices (US General Services Administration 2018, 2020a); product and service categories (US General Services Administration 2020b); and geography files to locate buyers and vendors (US Census Bureau 2010a, 2010b, 2010c, 2016).

We use these modifications to compute two measures of contract execution performance that are commonly used in the literature: cost overruns and delays (e.g., Decarolis 2014; Kang and Miller 2021; Decarolis et al. 2020; Carril 2022). Because the data contain the total sum of payments and the completion date expected at the time of the award, we can construct measures of cost overruns and delays by comparing these expectations to the realized payments and duration.[15] A few considerations suggest that these are meaningful measures of performance. First, both overruns and delays are routinely collected for larger contracts and used to evaluate the execution of contractors.[16] Furthermore, our interviews with contracting officers confirmed that staying on budget and on time is an important priority for

[15] The FPDS data records whether the modifications are in or out of contract scope. Our analysis does not restrict to a specific type of renegotiation, although out-of-scope modifications are extremely uncommon in our sample.

[16] For example, the IT dashboard—which tracks the performance of large IT projects—scores projects based on a series of considerations, two of which are deviations with respect to budgeted cost and scheduled delivery. Similarly, the DOD is required to periodically report to Congress on the cost and schedule status of all Major Defense Acquisition Programs.

the buyer. Finally, Carril (2022) shows that these execution measures are positively correlated with contract quality assessments based on objective product and service characteristics, using data from a sample of large IT contracts for which quality is systematically measured (Liebman and Mahoney 2017).[17]

The analysis sample consists of all competitively awarded definitive contracts[18] (DCs) with award values between \$10,000 and \$40,000, awarded in fiscal years 2015 through 2019 by the DOD, for products and services other than R&D.[19] Supplemental Appendix Table B.1 presents summary statistics of the sample. In total, roughly 86,000 contracts have been awarded to almost 30,000 firms. These contracts are awarded by 597 contracting offices belonging to the Army, the Navy, or the Air Force. The expected contract duration is 54 days on average, and all contracts are awarded on a fixed-price basis. A noteworthy feature of this setting is that competition is limited: An average contract receives 3.5 offers, with one out of four contracts receiving a single offer.[20] Winning vendors are often geographically close to the contracting offices, with both located in the same state in two out of every three contracts. Of the suppliers, 75 percent are characterized as small businesses. One out of every four contracts experiences modifications during the execution stage, leading to 7.6 percent of average cost overruns (i.e., excess cost relative to the original award value).

We also observe rich information about the type of goods and services that are contracted upon. Each award is classified into one of 1,479 possible standardized four-digit alphanumeric codes. These can be aggregated into 101 broader two-digit product categories, 77 goods, and 24 services. Supplemental Appendix Table B.2 shows the top ten most common two-digit good and service categories. The most common product categories are ADP equipment/software, medical equipment and supplies, and maintenance and repair equipment.

II. The Effect of Competition on Contract Outcomes

This section studies the effects of publicizing procurement solicitations on contract outcomes. As described in Section IA, federal regulation introduces a publicity requirement at \$25,000. We interpret these requirements as an exogenous increase in the level of competition for the award: They increase the set of potential participants who are aware of the auction. This is the relevant measure of the competitive

[17] For a sample of FPDS contracts merged to IT projects in the IT dashboard, Carril (2022) shows that overruns and delays as computed here correlate positively and significantly with officers' evaluations, which need to be based on objective performance metrics such as, for example, "percent of the time that the system is available," "percent of servers reduced as a result of virtualization," "number of repeat customers using system," etc.

[18] Federal contracts can be broadly categorized into two types: DCs and indefinite delivery vehicles. 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 (2022) for more details. We simplify the analysis by focusing exclusively on DCs, which are well-defined requirements involving a bilateral relationship between a single government unit and a private firm.

[19] The DOD represents 58 percent of overall federal procurement spending and more than 60 percent in the restricted sample. We exclude R&D awards because they are subject to a unique set of acquisition rules; see FAR part 35.

[20] More than half of the awards are set aside for a particular type of firm (typically, small business). Set asides are a major factor in the acquisition strategy of the DOD, and contracting offices are required to meet specific set-aside goals. Even though they affect contract competition, we abstract away from that feature as we do not condition nor restrict our sample based on that margin. Importantly, set-aside requirements do not vary within the range of contract values that we study.

environment for bidders, as they place their bids without knowing the set of actual participants. We exploit the $25,000 threshold to provide evidence of the effects of promoting competition on contract award price and other contract outcomes. These results will be the basis for our model in Section III.

A. *Preliminaries*

For each contract in our data, we observe agencies' decisions to publicly solicit a contract in FBO before its award (a decision that we denote as $D \in \{0,1\}$). We leverage the variation introduced by the regulation, which discontinuously affects the likelihood of public solicitation at an arbitrary threshold $(\bar{p} = 25{,}000)$ depending on the contract's expected award price $(\tilde{p})$. We do not observe ex ante estimated prices $\tilde{p}$, but only ex post realized prices p, which entails two empirical challenges. First, contracting officers know the policy threshold, which may generate incentives to modify the purchase to make the ex ante estimate fall below $\bar{p}$. This behavior would result in bunching on the distribution of ex ante prices, generating excess contracts estimated to be at or slightly below $\tilde{p} = \bar{p}$. Second, since prices ex post may differ from prices ex ante, estimating effects at the discontinuity may be subject to measurement error biases. In our case, publicity may affect prices due to enhanced competition; thus, the error distribution may differ depending on the publicity status of the contract.

To tackle these empirical challenges, we propose a method that uses the distribution of observed awards p and publicizing decisions D to (nonparametrically) recover information about the distribution of $\tilde{p}$, the distribution of the effects of publicity on price, and the extent of "manipulation."[21] Intuitively, the method hinges on comparing the observed empirical distributions of award prices with estimated counterfactual distributions stripped of the confounding influence of bunching and competitive price effects. We use this framework to estimate price effects and correct (and bound) RDD estimates on nonprice outcomes accounting for the aforementioned confounds. Since these corrections ultimately have a modest effect on our final results, we leave most details about the method to Supplemental Appendix D.[22]

Section IIB discusses the price effects of publicity obtained from the density analysis approach. Section IIC describes the estimation of publicity effects on nonprice outcomes relying on corrected RDD methods. Sections IID, IIE, and IIF provide interpretation of the estimated effects.

[21] By manipulation we mean any decision ex ante that modifies the requirement with the sole purpose of arriving at a different price estimate. The term follows the literature on regression discontinuity, which refers to this as "manipulation of the running variable." However, it is noteworthy that this behavior need not involve any wrongdoing.

[22] The key intuition for why these corrections are unimportant in our case can be better understood using a standard "donut-RD" logic: Since both sources of measurement error would be most pronounced right around the discontinuity cutoff, then RDD estimates should be sensitive to the exclusion of small subsets of observations around the threshold. In contrast, we show below that baseline estimates are quite robust to excluding windows around $25,000.

B. *Analysis of Contract Price Distributions*

Let $p_t(\tilde{p}_t | D_t)$ be the potential log-price that we would have observed for contract t, as a function of ex ante estimates $\tilde{p}_t$ and a publicity decision of $D_t \in \{0,1\}$. We assume that the unconditional distribution of potential contract prices is smooth and that publicizing solicitations leads to a log-linear random price effect for contract t—that is, $p_t(\tilde{p}_t | D_t = 1) = \tilde{p}_t - \gamma_t$, with $\gamma_t \sim F_\gamma(\cdot)$.

We estimate $E[\gamma_t]$ from the observed distributions of publicized and nonpublicized contracts. The intuition of our method is based on three observations. First, relative to a counterfactual with no price effects of publicity (i.e., $\gamma_t = 0$ for all t), the observed distribution of publicized contracts should be shifted horizontally by $E[\gamma_t]$. Second, by definition, the distribution of nonpublicized contracts is not affected by publicity effects γ_t. Third, we expect the counterfactual price distribution of the total number of contracts (the sum of those with and without publicity) to be smooth around the discontinuity, even though the threshold regulation generates discontinuities on each conditional distribution. These three observations motivate our method. We pick a value for $\widehat{E[\gamma_t]}$ and "undo" the price effects of publicity by shifting the distribution of publicized contracts, which we then add to the nonpublicized contracts. The "right" value of $\widehat{E[\gamma_t]}$ will satisfy the smoothness of the overall distribution and an integration constraint.

Supplemental Appendix D shows how this logic can be extended to nonparametrically identify the full CDF of price effects γ_t given the observed distributions of realized prices conditional on publicity status, $f(p_t | D_t = 0)$ and $f(p_t | D_t = 1)$. Moreover, the analysis is robust to strategic bunching in the distribution of nonpublicized awards, and the extent of this behavior is also identified using similar arguments. The key is that strategic bunching affects only the distribution of nonpublicized awards so that price effects and bunching are separately identified from the two observed distributions.

Figure 1, panel A depicts the (nonparametric) estimate of the CDF of γ_t, along with a local polynomial smoothing. We find that publicity leads to an average reduction in award price of 0.06 log points (SE: 0.02), equivalent to \$1,456 at the discontinuity. The full distribution shows that publicizing contract opportunities reduces award prices for 83 percent of the contracts. Supplemental Appendix Table B.3 provides more details about the mean and variance of price effects and displays subgroup analyses. We find that price effects are higher for services, and the effects are larger for more complex contracts.[23]

Supplemental Appendix Figure A4 shows the density distributions of both publicized and nonpublicized contracts, stripped down from price effects and strategic bunching responses. From the distribution of nonpublicized awards (panel A), we can directly compute the excess bunching below the threshold, explained by agencies' desire to avoid the publicity mandate. We estimate that the excess mass right below the discontinuity equals 12 percent of the value of the density at the threshold. This magnitude will be used to account for the effects of this manipulation on our

[23] By complexity, we refer to the average cost overruns for all contracts in the product category valued under \$20,000. This is discussed later in the paper.

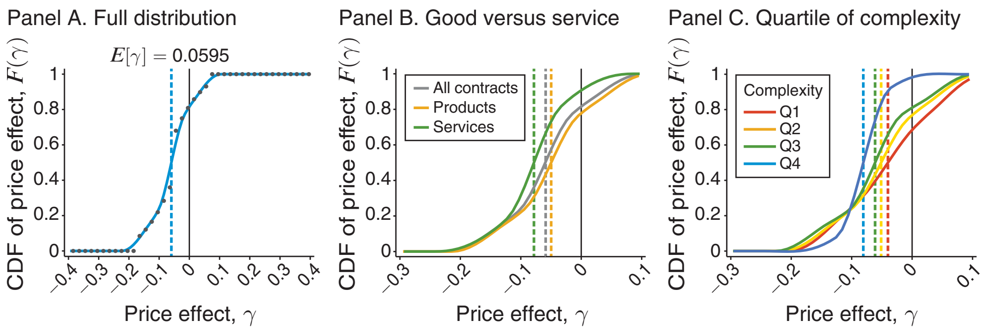

Figure 1. Distribution of Price Effects

Notes: This figure presents the estimated cumulative distribution function (CDF) of the price effects of publicity (γ). Panel A shows the cumulative CDF of all contracts in the sample. Every gray dot shows actual point estimates given a discretization of the support of γ. The blue line corresponds to a kernel fit. The dashed vertical line corresponds to the estimated mean effect. Panel B shows the CDF of price effects separating contracts for goods and services. Panel C shows the CDF of price effects by quartile of complexity.

RDD estimates in Section IIC. However, we can infer that since the extent of bunching is arguably modest, its impact on our estimates will also be limited.

Finally, in panel B, we compare the empirical distribution to the sharply discontinuous distribution of publicized awards that would be observed if $\gamma_t = 0$ for all t. It is evident how the distribution of γ_t smooths out the discontinuity in the density of publicized contracts. As noted by existing literature, observing the assignment variable with error biases the estimated effects toward zero in the RDD setting (Lee and Lemieux 2010; Davezies and Le Barbanchon 2017; Pei and Shen 2017). We leverage the estimated distribution of γ_t to correct for this factor in Section IIC.

C. *RDD: Estimating Effects on Nonprice Outcomes*

In this section, we leverage the discontinuous nature of the publicity requirements to gauge the effects of publicity on a set of other relevant outcomes, including the number of bids, characteristics of the winning bidder, and post-award contractor performance. We use the estimates of price effects and bunching to adjust the RDD estimates accounting for these factors.

Empirical Framework.—Consider specifications of the following form:

$$Y_t = \alpha + \tau \cdot D_t + f(\tilde{p}_i) + X_t'\zeta + \epsilon_t, \tag{1}$$

where the coefficient of interest is τ, the effect of publicizing a solicitation on contract outcome Y_t. In the standard RDD, we obtain an estimate of $\hat{\tau}_{IV}$ by instrumenting D_t with the discontinuity in publicity requirements. The first stage of this IV procedure is of the form

$$D_t = \kappa + \delta \cdot \mathbf{1}[\tilde{p}_t > \bar{p}] + h(\tilde{p}_t) + X_t'\eta + \nu_t, \tag{2}$$

for some smooth function $h(\cdot)$. 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—that is,

$$Y_t = \mu + \phi \cdot \mathbf{1}[\tilde{p}_t > \bar{p}] + g(\tilde{p}_t) + X_t'\psi + \xi_t. \tag{3}$$

Consider first a naive RDD, described by versions of equation (1), equation (2), and equation (3), where we simply replace ex ante prices $\tilde{p}_t$ by realized observed prices p_t. The estimates obtained from this analysis will be identical to the true RDD if there are neither price effects ($\gamma_t = 0$ for all t) nor bunching responses. The larger these effects are, the more the estimates from the naive RDD will differ from the true parameters. Given this, we take the naive RDD as our baseline and sequentially implement corrections to account for price effects and bunching responses, showing how these elements affect the estimation transparently.

In Supplemental Appendix D.3, we describe in detail the first of such corrections, namely a method to recover the causal parameters of interest in the presence of price effects γ_t. The key result is that, under our modeling assumptions, we can write the conditional expectation of contract outcomes given observed prices $E[Y_t | p_t]$ as an explicit linear function of the causal parameters that we seek to recover, plus objects that we can directly observe or estimate. This function depends on observed prices p_t, observed treatment probabilities π_D, and moments of the distributions of price effects F_γ (which we obtained from the density analysis). We then use this result to estimate the causal parameters using OLS.[24]

On the other hand, we can account for the effect of bunching responses by following the results from Gerard, Rokkanen, and Rothe (2020). These authors derive sharp bounds on treatment effects for the RDD in the presence of bunching. The simple argument is that if one can estimate the extent of manipulation in the running variable, which in our case corresponds to the excess mass below the threshold among untreated units (nonpublicized contracts), then one can derive bounds on treatment effects by assuming that these units are the ones with either the highest or the lowest values of the outcome variable Y_t. Intuitively, these are computed under the "worst" and "best" case scenarios regarding how selection can influence RDD estimates. In Supplemental Appendix D.4, we explain in detail how to derive these bounds in our setting and calculate them using our estimate of excess bunching obtained in our density analysis.

Effects on Nonprice Outcomes:

Baseline RDD Results.—We start with our baseline results—which ignore the possible influence of price effects or bunching—and then sequentially apply corrections to account for the specific issues present in our setting. We estimate equation (1), equation (2), and equation (3), assuming that $\tilde{p}_t = p_t$. In our baseline specifications, we use a simple linear fit for $g(\cdot)$ and no controls X_t, but we also present results from the robust local polynomial approach proposed by Calonico,

[24] We also show in Supplemental Appendix D.3 that this logic can be easily extended to accommodate measurement error in ex ante prices, so that $\tilde{p}$ is only an unbiased but not necessarily perfect forecast of $p^0(\tilde{p})$.

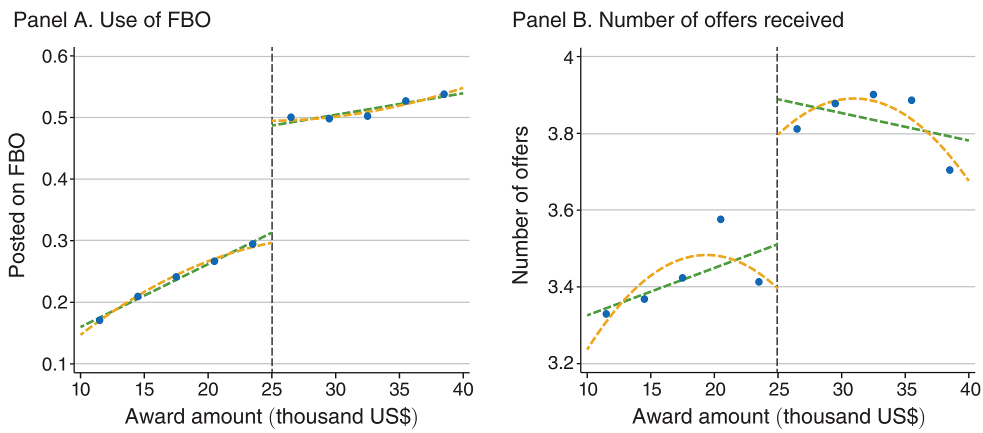

Figure 2. Publicizing Requirement and Intensity of Competition

Notes: Panel A and panel B respectively show the fraction of contracts posted on FBO and the number of offers received, as a function of award amounts. Blue dots represent average outcomes by bins of award amounts. Colored dashed lines represent linear and quadratic fits at each side of the $25,000 threshold. The data sources are FBO and FPDS-NG. The sample consists of competitive, non-R&D, DCs and purchase orders, with award values between $10,000 and $40,000, awarded by the DOD in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 length.

Cattaneo, and Titiunik (2014). We present these RDD results visually by plotting binned scatters of equation (2) and equation (3). In the next section, we explicitly assess how these baseline estimates change as we consider the impact of price effects and (or) bunching responses.[25]

The results for the first stage equation (2) are presented graphically in Figure 2, panel A. We see that the use of FBO jumps sharply past the $25,000 threshold of award amounts. The share of publicly solicited contracts in the government platform increases from roughly 30 percent at or slightly below $25,000 to 50 percent right above this threshold.

The reduced form specifications (equation 3) are estimated on three sets of outcomes: the number of bids received, winning vendor characteristics (including its relationship with the awarding office), and post-award performance. Most of the existing literature has studied these variables independently.[26] By studying them jointly, we can comprehensively understand the mechanisms and implications of policies oriented to enhance competition.

Figure 2, panel B shows how posting solicitations on FBO impacts the number of offers a contract receives around the threshold. Contracts right above $25,000 (which are more likely to be publicly solicited) receive roughly 0.4 more bids. The magnitude of the increase in the number of offers is considerable, given that the policy only changes the likelihood of a publicized solicitation by around 20 percentage points.

[25] Supplemental Appendix Figure A5 presents RDD plots for baseline variables. We find that baseline contract design characteristics are balanced around the threshold, except for a small difference in the share of goods versus services. All of our baseline estimates are robust to including a service dummy as a control.

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

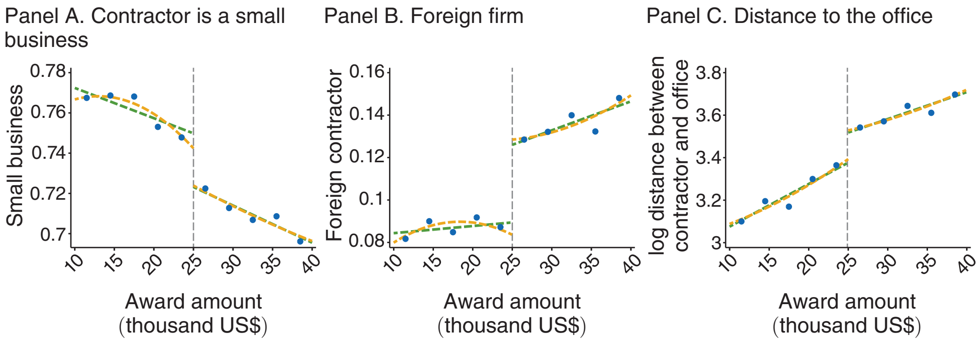

Figure 3. Publicity and the Characteristics of the Winning Firm

Notes: Panel A, panel B, and panel C respectively show the fraction of awarded contractors that are small businesses, the fraction of awarded contractors that are foreign, and the natural logarithm of the distance (in miles) between the contracting office's location and the contractor location, as a function of award amounts. Blue dots represent average outcomes by bins of award amounts. Colored dashed lines represent linear and quadratic fits at each side of the \$25,000 threshold. The data sources are FBO and the FPDS-NG. The sample consists of competitive, non-R&D, DCs and purchase orders, with award values between \$10,000 and \$40,000, awarded by the DOD in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of \$3,000 length.

These results indicate that expanding the set of potential bidders through the public posting of solicitations leads to the desired goal of increasing the number of actual participant bidders. However, it does not necessarily imply that these new offers affect the equilibrium allocation of the contract since new marginal bidders may not be competitive. Figure 3 shows that this is not the case. In panel A, we see that publicized contracts are awarded to vendors that are relatively larger, as measured by a reduction of the probability of awarding the contract to a small firm.[27] This "penalty" for small businesses is interesting because it goes against the stated goals of the publicity regulation (FAR part 5). Panel B and panel C show that publicized contracts are more likely to be awarded to foreign firms or firms that are geographically more distant from the contracting office location. These results suggest that marginal entrants attracted by public solicitation do win awards with a positive probability.

To measure the impact on post-award contract performance, we use two measures that are commonly used in the literature: cost overruns and delays. We compute these as the difference between the ex post realized sum of payments and duration of the project, and the expected value of these variables at the time of the award. Figure 4 presents the results. We find that the share of contracts with overruns and the share of contracts with delays increase by 2 and 1.5 pp, respectively. These differences are statistically and economically significant, considering the magnitude of the first stage. These results show that the execution of publicized contracts tends to result in poorer performance outcomes, including ex post costs. Supplemental Appendix Figure A6 shows effects on additional performance-related variables: the number of post-award contract modifications, cost-overrun dollars as a share of the original award, and days of delay relative to the expected schedule. These results

[27] The Small Business Administration defines size standards by NAICS industry. These standards depend on the number of employees and/or annual revenue. As a reference, the revenue standard for building cleaning services (NAICS code 561720), a common category in the sample, is \$19.5 million per year.

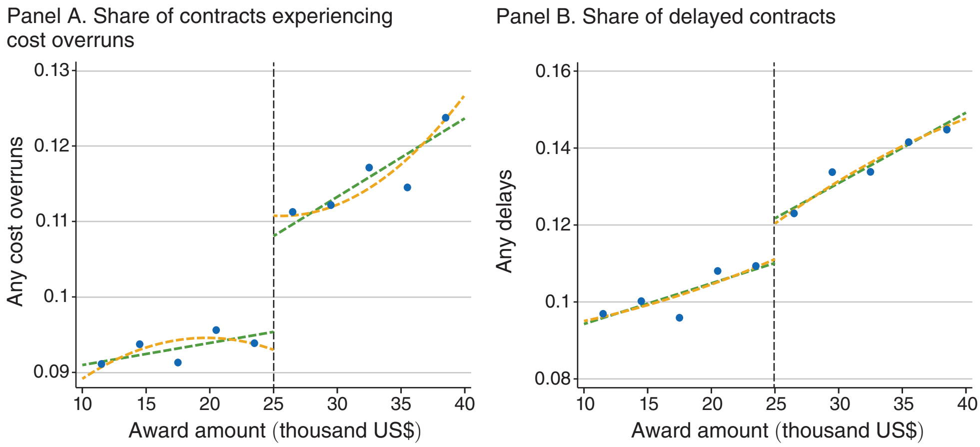

Figure 4. Publicity and Post-Award Contract Performance

Notes: Panel A and panel B respectively show the fraction of contracts experiencing cost-overruns and the fraction of contracts experiencing delays, as a function of award amounts. Cost overruns are measured as the difference between total obligated dollars and obligated dollars at the time of the award. Delays are measured as the difference between total duration of the contract and expected duration at the time of the award. Blue dots represent average outcomes by bins of award amounts. Colored dashed lines represent linear and quadratic fits at each side of the $25,000 threshold. The data sources are FBO and the FPDS-NG. The sample consists of competitive, non-R&D, DCs and purchase orders, with award values between $10,000 and $40,000, awarded by the DOD in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 length.

align with the findings presented in Figure 4: Publicized contracts experience more problems during the execution stage. Supplemental Appendix Figures A9–A16 illustrate how these effects vary by agency and type of purchased product.

Adjusted RDD Results.—In this section, we present a series of refinements to our baseline RDD results. First, we explore the robustness of our baseline linear specification with the estimator proposed by Calonico, Cattaneo, and Titiunik (2014), which uses robust local polynomial fits. Second, building upon the results of our density analysis in Section IIB and further explained in Supplemental Appendix D, we adjust the baseline RDD estimates to account for the observed running variable (award price) being subject to both treatment effects (price effects of publicity) and potential manipulation (bunching).

Table 1 presents reduced-form estimates for each relevant outcome variable. The first column shows the coefficient of our naive linear RDD using ordinary least squares (OLS). These results replicate the RDD plots discussed earlier. Column 2 presents Calonico, Cattaneo, and Titiunik's (2014) local polynomial estimates with robust bias-corrected standard errors. Overall, nonlinear estimates are similar in magnitude and significance to simple OLS estimates. The third column presents estimates that account for price effects in the treatment group (i.e., publicized contracts), following the method explained in Supplemental Appendix D.3. The correction for price effects is relatively modest and, in most cases, tends to amplify the baseline results. This is consistent with the fact that the price effects smooth out the discontinuity for the treatment group: Under naive estimation, some publicized contracts are observed below the threshold when their original (ex ante) price was above it.

Table 1—Reduced-Form RDD Estimates and Corrections

Dependent variable	OLS (1)	CCT (2)	Price effect adjustment (3)	Manipulation bounds (4)	Price effect + manipulation bounds (5)
Number of offers	0.3569	0.5447	0.3526	[0.2766, 0.5344]	[0.3075, 0.4506]
	(0.0677)	(0.1052)			
One offer	−0.0191	−0.0235	−0.0204	[−0.0274, 0.0052]	[−0.0249, −0.0070]
	(0.0064)	(0.0108)			
Log distance firm–office	0.1392	0.1199	0.1909	[0.0289, 0.2688]	[0.1303, 0.2619]
	(0.0481)	(0.0817)			
Foreign firm	0.0357	0.0508	0.0375	[0.0328, 0.0519]	[0.0358, 0.0465]
	(0.0045)	(0.0078)			
New firm	0.0348	0.0371	0.0427	[0.0186, 0.0518]	[0.0341, 0.0515]
	(0.0081)	(0.0136)			
Small business	−0.0277	−0.0295	−0.0265	[−0.0523, −0.0195]	[−0.0399, −0.0219]
	(0.0065)	(0.0110)			
Any cost overrun	0.0135	0.0246	0.0144	[0.0103, 0.0263]	[0.0127, 0.0216]
	(0.0045)	(0.0077)			
Cost overruns (relative dollars)	0.0095	0.0161	0.0127	[0.0053, 0.0179]	[0.0103, 0.0174]
	(0.0058)	(0.0100)			
Any delay	0.0130	0.0151	0.0143	[0.0093, 0.0271]	[0.0123, 0.0222]
	(0.0047)	(0.0081)			
Delays (days)	2.3262	4.0361	2.7491	[1.2282, 5.3550]	[2.1499, 4.4653]
	(2.0388)	(3.4952)			
Number of modifications	0.0375	0.0619	0.0395	[0.0204, 0.0926]	[0.0300, 0.0701]
	(0.0173)	(0.0301)			

Notes: This table shows RDD estimates of the reduced-form relationship between a series of outcome variables and an indicator of whether a contract award price exceeds $25,000. Each estimate comes from a separate regression. Coefficients in column 1 use a linear fit above and below the discontinuity. Coefficients in column 2, "CCT," correspond to the robust local polynomial method proposed by Calonico, Cattaneo, and Titiunik (2014). Column 3 applies a correction to the estimates in column 1, accounting for the existence of price effects, following the method proposed in Supplemental Appendix D.3. Column 4 shows bounds on the reduced-form coefficient in column 1, accounting for the possibility of "running variable manipulation" (i.e., bunching), following the method proposed in Supplemental Appendix D.4. Column 5 shows bounds on the adjusted reduced-form coefficient in column 4, accounting for both the existence of price effects and the possibility of "running variable manipulation" (i.e., bunching). Standard errors for the coefficients in columns 1 and 2 are shown in parentheses.

The next two columns present partial identification estimates that account for bunching responses. Column 4 shows lower and upper bounds without accounting for price effects, while the fifth column shows bounds that adjust for price effects. Notably, since the magnitude of bunching is modest in our context, the bounds presented are relatively narrow, which tells us that bunching does not pose a serious threat to the interpretation of our results. Interestingly, the lower bounds in column 5 tend to be very close to our baseline estimates. This implies that the downward bias introduced by price effects on the naive estimates of column 1 is of a similar magnitude to the worst-case upward bias introduced by bunching responses.

Taken together, these results imply that the strong visual evidence presented in Figure 2, panel A through Figure 4 is robust to the potentially confounding influence of price effects and strategic bunching by the buyer. Intuitively, the reason is that our estimates are not overly reliant on data that are very close to the discontinuity. As it is visually apparent from the figures and corroborated in Table 1, simple linear estimates that rely on a wider window of data are not too different from estimates that give a higher weight to the data very near the cutoff. This is important because the potential influence of price effects and strategic bunching would most strongly affect this area of the distribution. Hence, if these confounds were empirically relevant, the RDD estimates would be highly sensitive to excluding data points very close to $25,000. In contrast, Supplemental Appendix Table B.4 implements a series of "donut-RD" specifications (Barreca et al. 2011; Cattaneo and Titiunik 2022) and shows that our baseline linear estimates are robust to the exclusion of a window of varying width around the cutoff: For all window sizes considered, coefficients are statistically indistinguishable from the baseline.

D. *The Role of Contract Complexity*

Prior literature on incomplete contracts in procurement has emphasized the link between the underlying complexity of a transaction—an exogenous characteristic of the product being procured—and the level of costly ex post adaptation (Bajari and Tadelis 2001; Bajari, McMillan, and Tadelis 2009, Bajari, Houghton, and Tadelis 2014). Since the difficulty of specifying several possible contingencies varies across contracts, we should observe less variability in post-award performance if the purchase involves—for example, a standardized product rather than an ad hoc service. This may explain why some product categories in our data rarely experience execution issues ex post, while others present post-award modifications for most contracts. It may also imply that the effects on post-award performance that we documented above are heterogeneous across goods or services with different underlying complexity. Similarly, the impact of expanding competition on award prices is also likely to vary with complexity. For example, if bidders of relatively complex products are more heterogeneous in production costs, additional offers would lower contract award prices more than when contractors are homogeneous. Thus, the degree of contract complexity may shape how competition affects both prices ex ante and performance ex post.

To assess these mechanisms more directly, we leverage rich heterogeneity in our data, which features 1,918 distinct product categories. Using this information, we proxy the degree of complexity at the product category level based on the baseline level of execution performance. In particular, we define a category's degree of complexity as the average cost overruns experienced by all contracts in that category with an award below $20,000.[28] This leads to an intuitive classification, as shown in Supplemental Appendix Table B.5, which lists the complexity measure associated with the top and bottom product categories. Contracts for easy-to-specify purchases—like fuel, lumber, or medical supplies—receive the lowest complexity score. In contrast, contracts for more customized needs—for example, medical services, facility operation, and housekeeping—are associated with higher complexity.

[28] There are multiple ways of characterizing product complexity. We implemented different approaches, including using the standard deviation in performance, indexing multiple performance variables, and counting the number of words in the solicitation's description. These classifications lead to roughly the same rank of product categories and, thus, varying the definition does not threaten the general results. We present correlations between some of these measures in Supplemental Appendix Figure A7. We use the mean of cost overruns because it is transparent and easy to interpret. We get around the issue of classifying contract categories based on an outcome by focusing on awards below $20,000.

We divide the contracts in our sample into quartiles of complexity and reestimate both price effects and RDDs on performance, separately for each of the four groups. Supplemental Appendix Table B.3 shows estimates for the mean and standard deviation of price effects γ_t, separately for the full sample (column 1), goods versus services (columns 2 and 3), and each of the four quartiles of complexity (columns 4 through 7). Similarly, Figure 1 shows the CDFs of price effects for each of these groups. Although estimates become noisier as we divide the sample, we see suggestive evidence that large price effects are more concentrated among the most complex contracts. Our point estimates indicate that, on average, publicity reduces the prices of goods by 5 percent and services by 7.8 percent. This effect corresponds to 4 percent for the least complex quartile versus 9.6 percent for the top quartile of complexity.

The results are qualitatively similar for the impact of publicity on post-award performance. Figure 5 shows that the increase in overruns and delays that we reported in Figure 4 is driven by goods and services in the top quartile of complexity. We are unable to reject the null for the lower three quartiles.[29] Overall, it is noteworthy that both counteracting effects of competition—price reductions ex ante and overruns ex post—are more pronounced for complex contracts.[30] In Section III, we return to this trade-off and zoom in on the drivers of these effects.

E. *Evidence of Adverse Selection*

Our results show that increasing the pool of potential bidders through publicity generates changes to contract prices and subsequent contract execution. There are two classes of explanations through which we can rationalize the connection between publicity and contract performance: within-contractor and cross-contractor changes in performance. The within-contractor variation would imply that the same firm may perform systematically differently depending on the publicity status of the contract. This could stem from contractors' strategic choices (i.e., moral hazard) or features of the production technology. The cross-contractor variation would imply that publicity allows the participation of suppliers that are "different" and that their performance ability is unrelated to the contract's advertising (i.e., adverse selection).

To elucidate between these mechanisms, we leverage the fact that buyers often purchase the same product categories repeatedly over time, allowing us to observe multiple contracts for the same buyer-product combination, with variation in the size of the award and other characteristics of the contract. Moreover, on the supply side, we observe most contractors executing more than one contract for one or more different buyers. This variation allows us to test how much of the observed variation is due to contractors' "types," relative to variation "within" contractor.

[29] Supplemental Appendix Figure A16 shows regression discontinuity plots for cost overruns separating goods and services. Note that cost overruns increase for both types of contracts. However, both the baseline level and the magnitude of the jump are substantially larger for services.

[30] Because the publicity effects on award price and overruns are estimated using different methodologies—namely, density analysis and RDD—it is not straightforward to combine the two. This is why we leave a thorough analysis of the effects of publicity on the final price (i.e., award price plus overruns) for the model in Section III. Having said this, the evidence presented so far suggests: (i) adding the two effects implies an impact of publicity on final prices that is heterogeneous on complexity; (ii) given the estimated null effects on overruns for the least complex products, publicity likely reduces the final price for that group; and (iii) the trade-off is more nuanced for more complex contracts.

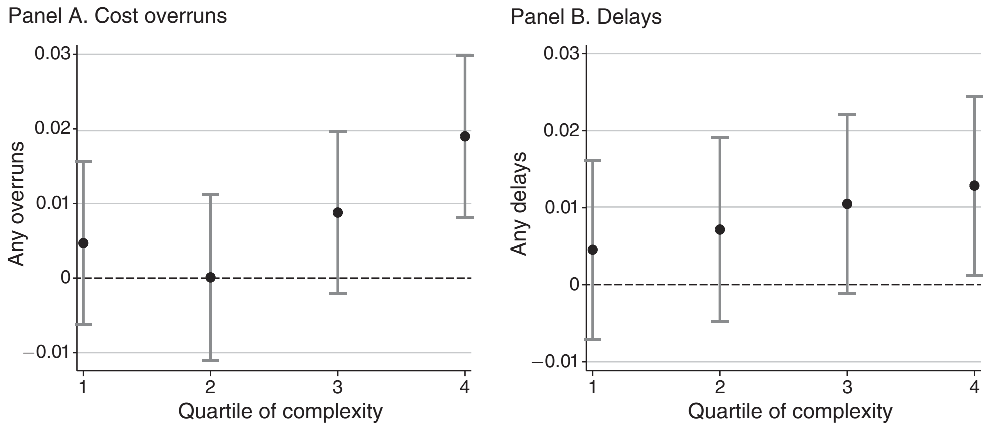

FIGURE 5. EFFECTS OF PUBLICITY ON POST-AWARD PERFORMANCE BY DEGREE OF COMPLEXITY

Notes: This figure shows four regression coefficients and their 95 percent confidence intervals. Each coefficient is an estimate of a RDD reduced-form coefficient in equation (3), per subgroup, estimated using (interacted) OLS. The dependent variable in panel A and panel B are indicators for any positive cost overruns and delays, respectively. The subgroups are determined by 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 \$20,000 in that category. The data source is the FPDS-NG. The sample consists of non-R&D DCs and purchase orders, with award values between \$10,000 and \$40,000, awarded by the DOD in fiscal years 2015 through 2019.

Table 2 presents the results of this exercise, where we reestimate RDD specifications assessing the effects of including contractor fixed effects. To make the different coefficients comparable, we constraint the sample in all columns to the set of firms for which we can estimate the fixed-effects specifications.[31] Columns 1 and 3 display the baseline IV estimates of the changes in performance induced by publicity, showing that the share of contracts experiencing overruns and delays jumps by 7.6 and 5.0 pp, respectively. Columns 2 and 4 show that incorporating contractor fixed effects shrinks the absolute value of the estimates substantially (to 2.2 and -1.3 pp, respectively), making both of them statistically indistinguishable from zero. Importantly, while there is no effect of publicity on performance when we look within contractors, columns 5 and 6 show that the actual competitive environment is indeed changing when we cross the threshold: The IV estimate of publicity on the number of offers barely changes with the introduction of contractor fixed effects, moving from 2.3 to 2.1 additional offers, both statistically different from zero and statistically indistinguishable from each other.

This evidence implies that most of the effects of publicity on contract performance are explained by variation across contractors, as opposed to within contractors. Conditional on the selected firm, performance does not significantly change when we cross the threshold, even though competition does increase substantially. We interpret these results as strong evidence in favor of adverse selection (i.e., that

[31] This leaves us with 81 percent of the contracts in the full sample (number of observations is 69,296, whereas Supplemental Appendix Table B.1 reports 85,661 contracts) and 41 percent of the firms (12,240 firms out of 29,641 firms reported in Supplemental Appendix Table B.1).

TABLE 2—IV-REGRESSION DISCONTINUITY ESTIMATES CONTROLLING FOR FIRM FIXED-EFFECTS

	Any cost overrun		Any delay		Number of offers	
	(1)	(2)	(3)	(4)	(5)	(6)
Estimate	0.076	0.022	0.050	−0.013	2.312	2.091
SE	(0.027)	(0.029)	(0.028)	(0.032)	(0.436)	(0.470)
Fixed effects?	No	Yes	No	Yes	No	Yes
Number of observations	69,296	69,296	69,296	69,296	69,296	69,296

Notes: This table shows instrumental variable estimates of the effect of appearing in FBO on contract outcomes using an RDD. The instrument corresponds to a dummy indicating whether the award value exceeds $25,000. The dependent variable in columns 1 and 2 is an indicator of having any positive cost overruns. The dependent variable in columns 3 and 4 is an indicator of having any positive delays. The dependent variable in columns 5 and 6 is the number of offers received. Columns 2, 4, and 6 include firm fixed effects. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the award time. Delays are computed as the difference between the actual duration of the contract and the expected duration at the award time. The data source is the FPDS-NG. The full sample consists of non-R&D DCs and purchase orders, with award values between $10,000 and $40,000, awarded by the DOD in fiscal years 2015 through 2019.

publicity brings in different vendors) as opposed to publicity inducing the same firms to change their behavior.

F. *Discussion*

To summarize, promoting vendor participation through publicity increases contract competition, as the average number of offers received rises substantially. The added competition translates into reductions in contract prices, but we also find that publicized contracts result in more cost overruns and delays. Our results show that promoting contract competition for (at least partially) incomplete contracts involves a trade-off: It reduces contract award prices at the cost of exacerbating adverse selection, leading to contractors characterized by lower execution performance. Furthermore, this trade-off is heterogeneous, with both price and performance effects depending on the degree of contract complexity.

While this policy analysis is informative of the effects of promoting competition on contract outcomes, it also presents some limitations. First, the estimated effects are local to the policy threshold of $25,000, so they may not be informative of the impacts for the rest of the sample. Second, they do not provide a clear representation of the mechanisms by which buyers' and sellers' characteristics and behavior shape market outcomes. Finally, our reduced-form analysis does not allow us to evaluate equilibrium conditions under alternative policy designs. To complement the previous analysis, we present and estimate a model of public procurement competition.

III. A Model of Competition Promotion, and Firms' Participation and Bidding Decisions

We develop and estimate an equilibrium model of publicity selection, firm participation, and bidding decisions in our public procurement setting. The ultimate goal is to estimate the model's primitives and study the implications of policy counterfactuals. We make modeling assumptions based on the setting's key features, aiming to transition from a theoretical model to an empirical one that can be estimated using the data available. Furthermore, we leverage the same variation highlighted in Section II—namely, the discontinuous nature of the publicity requirement—along with additional variation in market structure, to identify the model's parameters. Section IIIA introduces the theoretical model and discusses the auction's entry and bidding equilibrium strategies. Section IIIB describes the empirical implementation of the model, and in Section IIIC we discuss identification. The estimation approach is presented in Section IIID, and results are discussed in Section IIIE.

A. *Model*

A buyer offers a single and indivisible contract to potential contractors. Each potential contractor j must incur an entry cost $\omega_j \geq 0$ to learn her private cost to complete the task $c_j \in [\underline{c}, \bar{c}] \subset \mathbb{R}_+$ and bid for the contract. Both ω_j and c_j are assumed to be independent random draws. We model the potential bidders' choices in two stages. First, knowing the set of potential competitors $\mathbf{N}$, each potential bidder decides whether to incur the entry cost. After the entry stage, the $\mathbf{n} \subseteq \mathbf{N}$ firms that incurred entry costs privately learn their costs of completing the project and submit their bids b_j. The awarding mechanism is a first-price sealed-bid auction. The winner of the auction becomes the contractor, and its execution performance, q_j, is publicly observed once the contract is awarded.

Our analysis considers asymmetry between potential contractors: There are two types of firms, locals L and nonlocals NL. These firms differ in their production technology, which is characterized by the distribution of entry costs $H^k(\omega)$ and the joint distribution of project completion costs and execution performance $F^k(c,q) \equiv F^k_c(c) \cdot F^k_{q|c}(q\,|\,c)$, where $k(j) \in \{L, NL\}$ denotes bidder j's type. We assume that entry and project completion costs are private information and are distributed independently and identically within a type.

Contract Execution: Cost Overruns.—Throughout our analysis, we measure contractors' performance by the magnitude of cost overruns, which correspond to ex post realizations of unbudgeted costs. This variable is convenient as it can be directly benchmarked against the contract's awarded dollar value. Of course, this convenience comes at the cost of abstracting away from other (context-specific) execution costs.[32]

We assume that contractors draw q_j at the execution stage and fully pass through this cost shock to the buyer, leaving the utility of the contractor unchanged. Thus, potential differences in execution performance across contractors are explained by differences in production technologies that we take as given (i.e., types), rather than by strategic aspects. By estimating the primitive joint distributions $F^k(c,q)$ that capture these underlying production technologies, we will allow the data to inform us

[32] For example, Lewis and Bajari (2011) study the welfare gains associated with reducing delays in highway construction.

about the relationship between costs and performance.[33] This modeling choice is motivated by the evidence presented in Section IIE.

Publicity Choice.—The contract publicity status determines the set of potential participants as follows: If the contract solicitation is openly publicized, both local and nonlocal contractors learn about the contract opportunity; conversely, if the contract is not advertised, only local contractors receive the information. Our model allows the buyer to endogenously determine the set of potential contractors through the publicity choice, which is made taking into consideration (i) the likelihood that a contractor of each type is awarded the contract, (ii) the expected award price, and (iii) the expected execution performance in terms of cost overruns.

Equilibrium in the Bidding Stage.—Our analysis focuses on a type-symmetric equilibrium where bidders of type k follow the same bidding strategy $\beta_k(\cdot)$, mapping project cost c_j into a bid b_j. We assume that c_j is drawn independently from a type-specific continuous distribution $F_c^k(\cdot)$, with density $f_c^k(\cdot)$ and common support $[\underline{c}, \bar{c}] \subset \mathbb{R}_+$. The distributions of entry and production costs and the number of potential bidders of each type are common knowledge. Nevertheless, we assume that bidders do not observe the number of actual competitors of each type n_t^k, as in Krishna (2003) and Li and Zheng (2009).

Our setting considers two possible scenarios: with and without publicity. If the contract solicitation is publicized, then both local and nonlocal firms can participate. In this case, the expected payoff of bidder j, with type $k(j)$, cost realization c, and market composition $\mathbf{N} = \left(N^{k(j)}, N^{-k(j)}\right)$ is,

$$\Pi_j(b, c \mid \mathbf{N}) = (b - c)\,\mathbf{P}_j(b \mid \mathbf{N}), \tag{4}$$

where $\mathbf{P}_j(b \mid \mathbf{N})$ is the probability that bidder j wins the auction when bidding b. Given that the bidders don't observe the composition of actual competitors in the auction, the probability of winning depends on the entry probabilities of the other potential bidders:

$$\mathbf{P}_j(b \mid \mathbf{N}) = \left(\sum_{l=2}^{N^{k(j)}} \rho_l^{k(j)} \left(1 - G^{k(j)}(b)\right)^{l-1}\right)\left(\sum_{l'=1}^{N^{-k(j)}} \rho_{l'}^{-k(j)} \left(1 - G^{-k(j)}(b)\right)^{l'}\right), \tag{5}$$

where $G^k(b)$ is the distribution of equilibrium bids of type-k bidders, ρ_l^k is the probability that the number of actual bidders of type k is equal to l, and $-k$ denotes the other type of potential contractors. The optimal bidding is obtained from the usual first-order condition:[34]

$$(b - c)\,\mathbf{P}_j'(b \mid \mathbf{N}) + \mathbf{P}_j(b \mid \mathbf{N}) = 0. \tag{6}$$

[33] In Supplemental Appendix E, we further discuss the implications and limitations of our modeling assumptions.

[34] As noted by previous research on asymmetric auctions (Lebrun 1999; Bajari 2001; Maskin and Riley 2003a, b), the Lipschitz conditions are not satisfied in this case. Bidding strategies cannot be solved analytically; they require numerical methods. Campo, Perrigne, and Vuong (2003) and Brendstrup and Paarsch (2003) discuss nonparametric identification of cost functions in this setting.

If a contract solicitation is not publicized, only local firms can bid—that is, the number of potential nonlocal contractors is zero. In this case, the bidding problem is symmetric, as only one type is involved. Local suppliers observe the contracts' publicity status and, hence, the number of potential competitors.[35]

Equilibrium in the Entry Stage.—We focus on type-symmetric entry equilibrium strategies, φ^L, φ^{NL}. Firms compare the ex ante expected profit conditional on entry to their entry cost ω_j, where ω_j is independently drawn from a type-specific continuous distribution $H^k(\omega)$ with common support $[\underline{\omega}, \bar{\omega}]$. Firms with entry costs below their expected profit decide to incur the entry cost to learn about their cost of completing the project. The ex ante (expected) profits from participating are given by

$$(8)\quad \bar{\Pi}^k(\varphi^k, \varphi^{-k}) = \int_{\underline{c}}^{\bar{c}} \Bigg(\sum_{n^k-1, n^{-k} \subset N^k-1, N^{-k}} \pi^k(c | n^k - 1, n^{-k}; \mathbf{N}) Pr(n^k - 1, n^{-k} | \mathbf{N}) \Bigg) dF_c^k(c),$$

where $\pi^k(c \,|\, n^k - 1, n^{-k})$ denotes the expected profit of potential entrant with type k in an auction where the set of participants includes $n^k - 1$ and n^{-k}, and participants use equilibrium bid strategies $\beta_k(\cdot)$ and $\beta_{-k}(\cdot)$, respectively. The terms φ^k and φ^{-k} are the entry probabilities of each type. Because entry decisions are made simultaneously, the equilibrium condition is characterized by a type-specific entry cost threshold $\bar{\omega}_k$, such that firms whose entry cost is below their type-specific threshold participate.[36] Finally, when the contract is not publicized, only locals can participate, and thus the participation problem becomes symmetric. Therefore, for a given contract, the local type's participation threshold differs depending on whether the contract was publicized.

B. *Empirical Model*

Based on the equilibrium conditions of the general model, we proceed to describe its implementation in the empirical setting. A contract solicitation t is characterized by $(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t)$, $\mathbf{x}_t$ corresponds to a set of baseline characteristics of the contract, $\mathbf{z}_t$ are entry-cost shifters, and u_t captures unobserved project heterogeneity reflecting project attributes not included in the data but that impact firms' bidding behavior. In the empirical implementation below, $\mathbf{x}_t$ includes the contract's good or service category, associated complexity measure, expected duration, and location.

[35] Thus, the probability of winning a contract that is not publicized is

$$(7)\quad \mathbf{P}_j(b \,|\, \mathbf{N}) = \left(\sum_{l=2}^{N^{k(j)}} \rho_l^{k(j)} \left(1 - G^{k(j)}(b)\right)^{l-1} \right).$$

[36] In equilibrium, the entry probabilities are defined by the system of equations:

$$\varphi^L = H^L\left[\bar{\omega}_L(\varphi^L, \varphi^{NL})\right]$$

$$\varphi^{NL} = H^{NL}\left[\bar{\omega}_{NL}(\varphi^L, \varphi^{NL})\right]$$

type-specific equilibria exist by Brouwer's fixed point theorem. We numerically verified the uniqueness of the equilibrium entry probabilities within our estimation routine (Krasnokutskaya 2011; Roberts 2013). Espin-Sanchez et al. (2021) discuss sufficient conditions for equilibrium uniqueness in entry games with private information.

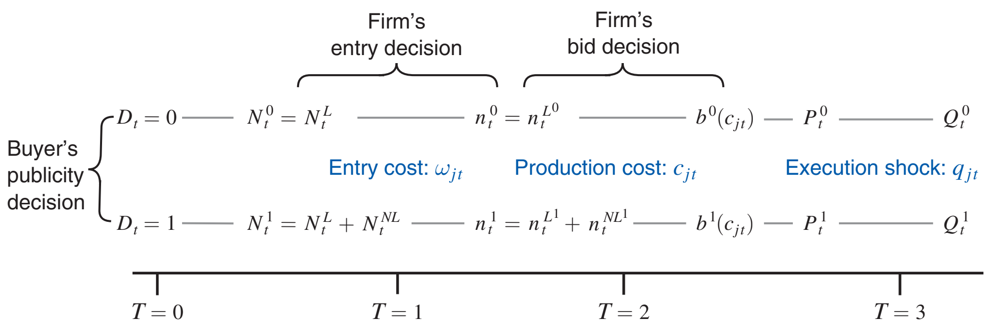

Figure 6. Timing of the Model

On the other hand, for $\mathbf{z}_t$, we use variables related to the timing of the solicitation, particularly whether it was awarded at the end of the fiscal year.[37] Finally, $\mathbf{N}_t = \left(N_t^L, N_t^{NL}\right)$ denotes the number of potential contractors of each type.[38] The model proceeds in four stages, depicted in Figure 6. The stages are as follows:

$T = 0$: Publicity Decision.—The buyer observes $\left(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t\right)$ and decides whether to publicize the contract, $D_t \in \{0, 1\}$, in order to maximize expected utility. Contract publicity status determines the set of potential bidders.

$T = 1$: Entry Decision.—Each firm that learns about the contract opportunity observes $\left(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t\right)$. They draw individual and private realizations of entry costs and simultaneously decide whether to participate.

$T = 2$: Bid Decision.—Active bidders draw a realization of the production cost and decide the magnitude of their bid. The contract price equals the lowest bid submitted.

$T = 3$: Execution Stage.—Once the contract is finalized, the performance of the selected contractor is observed and corresponds to the realization of an execution shock in terms of cost overruns.

[37] As documented by Liebman and Mahoney (2017), the volume of contracting activity at the end of the fiscal year is disproportionately high, which may reduce the attention that any given vendor can devote to each solicitation.

[38] Identifying the potential number of bidders is not trivial (Athey, Levin, and Seira 2011; Krasnokutskaya and Seim 2011; MacKay 2022). We combine two methodologies. First, using the procedure described in Supplemental Appendix Section G.1, we classify and count the suppliers that ever won a contract for every buyer-product combination. The second method considers the maximum number of actual bidders for buyer-product auctions. This method is discussed by Athey, Levin, and Seira (2011) and Roberts (2013). It is rooted in the theoretical idea that if all potential bidders decide whether to enter simultaneously, with enough observations, the maximum number of observed bidders across observations will be equal to the total number of potential bidders. The maximum number of bidders of auctions that weren't publicized informs about the number of potential local bidders. In contrast, the maximum number of bidders of advertised contracts approximates the sum of local and nonlocal potential bidders. We define the number of potential bidders for every buyer product as the maximum of both approaches. Combining these two methods alleviates each of their potential weaknesses. The median number of potential local and nonlocal bidders is six and three, respectively.

Specification:

Publicity Decision.—We assume that the buyer is risk neutral, forms unbiased beliefs, and derives utility from expected contract outcomes log linearly:

$$U_t^D = U\left(\bar{P}_t^D, \bar{Q}_t^D, \bar{L}_t^D\right) = \lambda^P \bar{P}_t^D + \lambda^Q \bar{Q}_t^D + \lambda^L \bar{L}_t^D + \varepsilon_t^D, \tag{9}$$

where $\bar{P}_t^D$, $\bar{Q}_t^D$, and $\bar{L}_t^D$ are, respectively, the expectation of the log awarding price, log cost overruns, and likelihood that a local contractor wins. These expected outcomes are objects that depend on realizations of $(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t)$ and the publicity status $D_t \in \{0,1\}$. The terms $\bar{P}_t^D$ and $\bar{Q}_t^D$ are in log-dollar units, while $\bar{L}_t^D$ is a probability. The parameters λ^P and λ^Q capture standard price sensitivity. λ^L captures any form of favoritism unrelated to award price or execution performance.[39] Finally, ε_t^D is an idiosyncratic utility shock.

The publicity regulation kicks in when the expected award price without publicity is higher than \$25,000. This introduces a utility shift, η, which translates into a discrete jump in the probability of advertisement at the threshold. Intuitively, η captures the intensity of regulation enforcement above the threshold:[40]

$$D_t = 1 \Leftrightarrow U\left(\bar{P}_t^1, \bar{Q}_t^1, \bar{L}_t^1\right) + \eta \mathbf{1}\left(\bar{P}_t^0 > \log(25)\right) \geq U\left(\bar{P}_t^0, \bar{Q}_t^0, \bar{L}_t^0\right). \tag{10}$$

Entry and Bidding Decision.—Bidder j's execution cost for contract t is multiplicative: $c_{jt} = \tilde{c}_{jt} \cdot u_t$, where $\tilde{c}_{jt}$ is a firm-specific cost component that is private information of firm j, and u_t represents a common cost component that is known to all bidders but is unobserved by the researcher (Haile and Kitamura 2019). The distribution of the firm-specific cost component for type-k firms is given by $F_{\tilde{c}}^k(\cdot \mid \mathbf{x}_t)$, and each draw $\tilde{c}_{jt}$ is independent across j and t conditional on observables. The unobserved project heterogeneity is given by $u_t \sim K(\cdot)$, and is independent from project characteristics $(\mathbf{x}_t)$, entry-cost shifters $(\mathbf{z}_t)$, and the set of potential bidders $(\mathbf{N}_t)$.

We assume that bidders are risk neutral. Thus, the Bayes-Nash equilibrium bid function for type k is multiplicative: $\beta_k(c_{jt} \mid \mathbf{x}_t, u_t, \mathbf{N}_t) = u_t \cdot \tilde{\beta}_k(\tilde{c}_{jt} \mid \mathbf{x}_t, \mathbf{N}_t)$.[41] Each bidder submits a bid of $b_{jt} = \tilde{b}_{jt} \cdot u_t$, where $\tilde{b}_{jt} = \tilde{\beta}_k(\tilde{c}_{jt} \mid \mathbf{x}_t, \mathbf{N}_t)$ represents the bid for bidder j when u_t is one. Therefore, $\log(b_{jt}) = \log(\tilde{b}_{jt}) + \log(u_t)$, and the log

[39] We remain agnostic about the specific nature of this preference—which could reflect anything from convenience to path dependence to corruption—and focus on the impacts on total procurement spending.

[40] Since we have documented that contract "manipulation" (i.e., strategic bunching of price estimates) is unimportant in this setting, for simplicity, we abstract away from this action in the model.

[41] This formulation is discussed by Krasnokutskaya (2011), proposition 1. The author shows that when the cost function is multiplicative with unobserved heterogeneity, Bayes-Nash equilibrium bidding strategies are also multiplicative.

of the unobserved heterogeneity component acts as an additive mean shifter to the conditional distribution of log bids.[42,43]

Finally, we assume that entry costs ω_{jt} are independent across j and t conditional on observed project characteristics $\mathbf{x}_t$. In equilibrium, firms' participation behavior is characterized by type-specific thresholds, $\bar{\omega}_t^{k*}$. Thus, the number of actual bidders n_t^k from type $k \in \{L, NL\}$ distributes binomial with an individual entry probability of $\varphi^{k*}(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t)$ and N_t^k trials, where $\varphi^{k*}(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t) = H^k\Big(\bar{\omega}^{k*}(\mathbf{x}_t, \mathbf{z}_t, u_t, \mathbf{N}_t)\Big)$. Our model considers entry shifters $\mathbf{z}_t$, which capture market-level conditions that affect entry decisions.[44]

Contract Execution.—Contract execution is observed ex post and corresponds to the magnitude of cost overruns, drawn from the distribution $F_{q|c}^k(\cdot \mid \mathbf{x}_t)$. Thus, $F_{q|c}^k(\cdot \mid \mathbf{x}_t)$ is type dependent, varies by observables $\mathbf{x}_t$, and is directly observed in the data.

Equilibrium is characterized by the buyer choosing a contract's publicity status that maximizes her expected utility, and informed potential contractors entering and bidding optimally if expected profits exceed their entry costs. Finally, contract execution is revealed once the contract concludes.

C. *Identification*

We aim to identify the type-specific distributions of ω_{jt}, $\tilde{c}_{jt}$ and q_{jt}, the distribution of the unobserved heterogeneity u_t, and the parameters that govern the buyer's utility function. For every contract t, we observe in the data a set of model *inputs*$(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ and *outputs*$(D_t, \mathbf{n}_t, P_t, Q_t, L_t)$. The model is identified based on three main assumptions:

- Contract and market characteristics $(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ are exogenous.
- Idiosyncratic components of entry cost shocks are (conditionally) independent from production cost and execution shocks—that is, $\omega_{jt} \perp (c_{jt}, q_{jt}) \mid \mathbf{x}_t$.
- Unobserved heterogeneity u_t is independent across t with $\mathbb{E}[u_t \mid \mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t] = \mathbb{E}[u_t] = 1$.

[42] We assume that the buyer sets a shadow reserve price at the ninety-ninth percentile of the cost function of local contractors. Establishing a shadow reserve price acknowledges that the buyer is not necessarily forced to accept unreasonably high bids and allows us to discard equilibria in which firms bid infinity due to the chance of being the only bidder. This way, the reserve price is scaled by x_t and u_t, and is almost never binding in compliance with necessary conditions for identification (Krasnokutskaya 2011). Related papers either rule out data from auctions with only one bidder or assume that in the case of a single bid, the buyer operates as a second bidder (Li and Zheng 2009; Athey et al. 2011; Krasnokutskaya and Seim 2011).

[43] Note that the buyer's utility function is linear on the log-winning bid, so an auction's $\log(u_t)$ is an additive term on both sides of equation (10), implying that it cancels out in the buyer's publicity choice.

[44] In the absence of entry shifters, variation in actual entrants would still be obtained from changes in potential bidders $\mathbf{N}_t$. So while the entry-cost shifters $\mathbf{z}_t$ are not strictly necessary for identification, they generate additional variation in the number of actual entrants $\mathbf{n}_t$ for a given vector of observables $\mathbf{x}_t$. This is useful since the pass-through from $\mathbf{N}_t$ to $\mathbf{n}_t$ may be small and potentially ambiguous due to counteracting "entry" and "competition" effects (Li and Zheng 2009). Now, note that $\mathbf{z}_t$ provides additional variation only if it is uncorrelated with production costs. In our estimation, $\mathbf{z}_t$ is a dummy that takes value one if the contract was solicited in the last month of the fiscal year. The reasoning is that the timing of the auction affects entry costs (e.g., due to short solicitation notice) but is independent of contract costs. The latter is equivalent to assuming that the distribution of project cost is constant over time.

Identification in our model involves pinning down primitive distributions of the two types of bidders (locals and nonlocals). In our setting, the contract's publicity status determines the composition of participating bidders. The data from nonpublicized contracts inform about the distributions of local contractors, while nonlocal contractors are only observed on publicized contracts. Identification of these two distributions requires that the sample of publicized and nonpublicized contracts be randomly selected. Yet, buyers' publicity choices are driven by observables, unobservables, and utility shocks. In the spirit of the RDD discussed above, we leverage the discrete nature of the publicity requirement to obtain quasiexperimental variation in publicity adoption and, thus, identify type-specific distributions separately. We now discuss identification more specifically for the different components of the model.[45]

Bidding.—The empirical challenge involves separately identifying $F_{\tilde{c}}(\tilde{c})$ from $K(u)$. The identification argument relates to MacKay (2022) and builds upon exogenous variation in the number of bidders, $\mathbf{n}_t$.[46] In our setting, bidders observe auction characteristics and the set of potential competitors, $\mathbf{N}_t$, but do not know the set of actual competitors, $\mathbf{n}_t$. Equilibrium bidding strategies depend on the information they have in hand: Auction bidders that form the same beliefs about the competitive environment would set the same bidding strategies. Thus, the number of actual competitors $(\mathbf{n}_t | \mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ would depend on realizations of (random) individual entry cost shocks.

Exogenous variation in the number of entrants allows for identifying $N - 1$ expected order statistics of the bidding distribution for each $(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ combination. Since u_t is assumed independent, one additional competitor under the same bidding strategy is equivalent to one additional draw from the distribution of normalized bids, $G^k(\tilde{b})$. Restrictions over expected order statistics approximate the quantiles of $G^k(\tilde{b})$, and if $N \to \infty$, $G^k(\tilde{b})$ is exactly identified. The underlying cost distribution $F^k_{\tilde{c}}(\tilde{c})$ is pinned down from the distribution of $G^k(\tilde{b})$(Guerre et al. 2000; Campo et al. 2003). See Supplemental Appendix F.1 for more identification details and proofs.

Entry.—Potential bidders set a threshold for realizations of entry costs. They pay ω_{jt} and enter auction t only if the realization ω_{jt} is smaller than the expected profit of participating in the auction—that is, $\omega_{jt} < \bar{\omega}^k_t$. Since the probability of participating enters into the expected utility function, which defines the cutoff, the (fixed-point) equilibrium entry cutoff is characterized by a type-specific entry probability

[45] In what follows, we omit the distinction between $D_t \in \{0,1\}$ because it is taken as given by the bidders. To ease notation, when a contract t is publicized, the set of bidders has two dimensions: that is, $\mathbf{N}_t = (N^L_t, N^{NL}_t)$ and $\mathbf{n}_t = (n^L_t, n^{NL}_t)$. If a contract is not publicized, only locals could participate: $\mathbf{N}_t = (N^L_t, 0)$ and $\mathbf{n}_t = (n^L_t, 0)$.

[46] Alternative strategies to identify models with unobserved heterogeneity involve either stringent assumptions on auction participation or observing the full distribution of bids. Compiani et al. (2020) assumes the number of active bidders can be characterized by an (equilibrium) reduced-form relation, $n_t = \eta(x_t, z_t, u_t, N_t)$ that is weakly increasing in u_t, thus a realization of n_t inform about (unobservable) realizations of u_t. Roberts (2013) provides a similar identification argument but leveraging variation in auctions' reserve price. Alternatively, Krasnokutskaya (2011) follows a measurement error approach and builds upon deconvolution methods to separately identify the distribution of unobserved heterogeneity and individual cost functions. The latter requires observing at least two bids per auction.

$\varphi_t^{k*} = H^k\left(\bar{\omega}_t^{k*}\right)$. Identifying $H^k\left(\cdot\,|\mathbf{x}_t\right)$ from the data entails three steps. First, we observe the realized set of potential bidders that decide to enter each auction t, which implies that, with enough observations per combination of $\left(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t\right)$, we can estimate $\varphi^{k*}\left(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t\right)$. Then, we use equation (8) to back out the expected utility of entering, conditional on $\left(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t\right)$. The final step leverages variation in $\left(\mathbf{z}_t, \mathbf{N}_t\right)$ to construct combinations of $\left(\varphi^{k*}, \bar{\omega}^{k*}\,|\,\mathbf{x}_t\right)$ that pin down $H^k(\omega)\left(\cdot\,|\,\mathbf{x}_t\right)$.

Execution.—Contract execution performance (cost overruns) is given by (conditionally) independent random shocks. Thus, the observed distribution of cost overruns directly reveals $F^k_{q|c}\left(\cdot\,|\,\mathbf{x}_t\right)$. The joint distribution of $F^k\left(\tilde{c}, q\,|\,\mathbf{x}_t\right)$ is obtained from $F^k_{q|\tilde{c}}\left(q\,|\,\mathbf{x}_t\right) \cdot F^k_{\tilde{c}}\left(\tilde{c}\,|\,\mathbf{x}_t\right)$.

Buyer's Preference Parameters.—The buyer's taste parameters for price, overruns, and local contractors are identified from variation in contract and market characteristics $\left(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t\right)$. In particular, variation in the set of potential bidders determines the effects of publicity on price and on having a local winning the auction. The degree of complexity of the transaction helps pin down the potential scope for overruns ex post. Intuitively, keeping other factors fixed, if a transaction involves a fully specified product, there will be no differences in performance ex post, which shuts down that factor in the decision. Finally, the disutility from avoiding the publicity regulation is identified from the publicity choices below and above the threshold. Following the logic of the RDD in Section II, the difference in publicity intensity below and above the threshold directly informs about the "bite" of the policy, which reveals its effects on buyers' utilities.

D. *Estimation Approach*

We make functional form assumptions to characterize equilibrium conditions of each stage and take the model to data:[47]

Bidding and Execution.—We specify normalized equilibrium bids as a linear function of covariates, $\mathbf{x}_t$, and the set of potential bidders, $\mathbf{N}_t$:

$$\log\left(\tilde{b}_{jt}\right) = \left[\mathbf{x}_t, \mathbf{N}_t\right]' \alpha^k + \varepsilon^{\tilde{b}}_{jt}.$$

We further assume that $\log\left(u_t\right)$ is normally distributed with mean zero and variance σ_u^2. The observed distribution of cost overruns is censored at zero to accommodate the fact that most contracts stay right on budget (Eun 2018): $Q_t^k = \max\left\{0, \log\left(q_t^k\right)\right\}$, where $log\left(q_{jt}\right)$ is a linear function of observables:

$$\log\left(q_{jt}\right) = \mathbf{x}_t' \gamma^k + \varepsilon^q_{jt}.$$

[47] Our parametric assumptions are linked to related literature (Krasnokutskaya and Seim 2011; Hong and Shum 2002; Porter and Zona 1993). Moreover, Krasnokutskaya (2011) indicates that the distribution of firm-specific components and unobserved heterogeneity closely resembles log normality. Overall, our results are not sensitive to adding additional covariates or variations to the functional form. Our data provide enough variation to identify these distributions independently from the specific functional form.

We specify the joint distribution of bid and execution shocks as a bivariate normal:

$$\begin{pmatrix} \varepsilon_{jt}^{\tilde{b}} \\ \varepsilon_{jt}^{q} \end{pmatrix} \sim N\begin{pmatrix} 0 & \sigma_{\tilde{b}}^2 & \rho\sigma_{\tilde{b}}\sigma_q \\ 0 ' & \rho\sigma_{\tilde{b}}\sigma_q & \sigma_q^2 \end{pmatrix},$$

where $\sigma_{\tilde{b}}^2 = \exp\{\mathbf{x}_t'\nu^k\}$, $\sigma_q^2 = \exp\{\mathbf{x}_t'\xi^k\}$, and ρ specifies the correlation between bidding and execution shocks. Intuitively, $\rho < 0$ implies that vendors with negative bid shocks (who bid low due to low cost draws) tend to experience more cost overruns ex post. We specify $\rho = \mathbf{x}_t'\iota^k$.

Entry.—The equilibrium entry choices are characterized by type-specific probabilities, $\varphi_t^k(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$. We assume $\varphi_t^k(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t) = \Phi\left([\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t]'\tau^k\right)$, where $\Phi(\cdot)$ denotes the cumulative distribution of the standard normal distribution, and $\mathbf{z}_t$ are entry-cost shifters. Since the entry decisions are independent and simultaneous, the number of participating bidders $n_t^k(\cdot)$ follows a binomial distribution with N_t^k independent draws with a probability of success $\varphi_t^k(\cdot)$.[48]

Publicity Choice.—We specify that the difference of buyers' utility shocks $(\varepsilon_t^0 - \varepsilon_t^1)$ in equation (9) distributes standard normal, so that $Pr(D_t = 1) = \Phi\left(\lambda^P\tilde{\bar{P}}_t + \lambda^Q\tilde{\bar{Q}}_t + \lambda^L\tilde{\bar{L}}_t + \eta\mathbf{1}\left(\bar{P}_t^0 > 25\right) + \mathbf{x}_t'\zeta\right)$, where $\left(\tilde{\bar{P}}_t, \tilde{\bar{Q}}_t, \tilde{\bar{L}}_t\right)$ are the change in expected outcomes associated with publicity, leaving no publicity as the omitted category.[49] We include agency fixed effects as well as $\mathbf{1}\left(\bar{P}_t^0 > 25\right)$ to indicate whether the expected price without publicity is above the regulation threshold.

Our specification provides flexibility to allow all distributions to differ across locals and nonlocals. Moreover, we interact our main covariates with indicators of nonlocal bidders.

Estimating Dataset.—The data used to estimate the model is the same one used in previous sections, except for one additional restriction. To classify local and nonlocal vendors, we require buyer-product combinations to appear at least four times in the full database between 2013 and 2019, at least one of which should appear in FBO. This restriction rules out products that are purchased less often. Supplemental Appendix Table B.6 compares summary statistics for the relevant variables between this selected sample and the full sample used in Section II. Overall, the model sample is broadly representative of the full sample in terms of observables. However, given that the selection involves the buyer contracting the same product multiple times, the selected sample includes contracts for categories that are, on average, less durable, with a slightly higher probability of being publicized in FBO (0.39 versus 0.3). Finally, and

[48] Related papers either assume parametric distributions for the entry costs, which paired with the expected utility of entering map into well-defined type-specific entry probabilities (Krasnokutskaya and Seim 2011; MacKay 2022); or make functional form assumptions on the entry probabilities, which combined with expected utilities allow for recovering entry costs (Athey et al. 2011, 2013; Compiani et al. 2020). We follow the latter approach.

[49] Our estimation does not restrict the set of values for parameters λ^P, λ^Q, and λ^L. However, in general, we may expect that buyers dislike paying higher prices or experiencing overruns, so we expect λ^P and λ^Q to be negative.

consistent with the rest of the analysis, we estimate the model using contracts around the regulation threshold—that is, between \$10 thousand and \$40 thousand.

Estimation Procedure.—Our empirical model yields predictions about equilibrium conditions for suppliers' participation, bidding, and ex post execution, with and without publicity. We also characterize the buyer's publicity decision. Our estimation strategy proceeds using the simulated method of moments (SMM) (Mcfadden 1989; Pakes and Pollard 1989). We choose a vector of parameters θ to generate simulated moments that closely resemble key moments from the data. Using the parametrized primitives discussed previously, we simulate four sets of moments: participation decisions, bidding strategies, cost overruns, and publicity decisions.

Our simulation procedure starts with a set of data inputs $(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ of size T. Then, from every observation, we generate S random draws of u_t. Finally, our setting contemplates that the buyer decides based on expectations, which are formed conditional on $(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$ and u_t, integrating over Monte Carlo simulated distributions of award price, overruns, and the likelihood of a local winning. Although computationally involved, this method is useful to circumvent integrating over potentially nonlinear functions and provides enough flexibility to match theoretical moment functions that cannot be evaluated directly.

Formally, denote the target m_n as a vector of moments from the data. The analogous moments generated by simulating observations are denoted by $m_s(\theta)$. This vector depends on the parameters $\theta \in \Theta \subset \mathbb{R}^P$. The estimator minimizes the standard distance metric:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \left(m_n - m_s(\theta)\right)' W_n \left(m_n - m_s(\theta)\right), \tag{11}$$

where W_n is the weighting matrix.[50]

The vector of parameters is $\theta = \left(\alpha^k, \nu^k, \tau^k, \gamma^k, \xi^k, \iota^k, \overrightarrow{\lambda}, \eta, \zeta, \sigma\right)$. We use three sets of target moments. The first set of moments is a vector of means and variances of the outcome variables and their interaction with the relevant covariates. The outcome variables are the auction price, the number of bidders, an indicator for a local winner, the magnitude of cost overruns, an indicator for any cost overrun, and publicity choices. The second set of moments consists of means of these same outcome variables conditional on partitions of the domain of contract prices and are estimated separately for goods and services. Finally, the third set of moments corresponds to a vector of normalized observation frequencies on the relevant window of contract prices. Stacking these three vectors, we obtain the vector m_n of 109 moments that we seek to match with the model. We use the stochastic optimization algorithm differential evolution (Storn and Price 1997) to perform the objective minimization.[51] The details of the estimation procedure are discussed in Supplemental Appendix G.

[50] The quasioptimal weighting matrix, W_n, is obtained previously from a set of candidate parameters (Gourieroux, Monfort, and Renault 1993)

[51] This algorithm performs a (parallel) direct search approach; it does not rely on gradient methods for minimizing possibly nonlinear and nondifferentiable continuous space functions.

E. *Estimation Results*

We estimate the model's parameters of publicity selection, entry, bidding, and execution, and combine these estimates with model equilibrium conditions to recover the primitive distribution of production and entry costs for locals and nonlocals. These estimates are inputs for the policy counterfactuals in Section IV.

Estimates.—To facilitate the interpretation of coefficients, panel A in Table 3 presents marginal effects. Although the coefficients are estimated jointly, the marginal effects are presented in three different columns for each decision (entry, bidding, and execution). Supplemental Appendix Table B7 displays the underlying coefficient estimates with their corresponding standard errors.[52] Panel B shows the estimates associated with standard deviations and correlations, and panel C describes buyers' preference parameters in terms of marginal effects.

Several findings are worth highlighting. First, bidders are less prone to participate if the contract involves a service or a relatively complex product. Thus, auctions for these types of products are less competitive. In line with the evidence presented in Section IIE, nonlocal contractors are more likely to participate than locals, and bidders reduce their probability of entering if they expect more competition (i.e., face a higher number of potential bidders). As expected, one additional potential nonlocal competitor has a higher effect on entry than a local one.

Second, bids from nonlocals are 1 percentage point lower than bids from locals, and more potential competitors reduce the equilibrium bids. Another relevant feature is that unobserved heterogeneity is important in our data. Most of the variation in bidding is explained by common factors instead of variation between bidders within an auction. The standard deviation of (log) unobserved heterogeneity is 4.3 times larger than the bids' standard deviation when $\log(u_t) = 0$.[53]

Third, the execution shock depends on the transacted product: The mean of log-overrun shocks is substantially higher for more complex products. In line with the reduced-form results, the difference in cost overruns between locals and nonlocals is sizable, with nonlocals having a mean shock of 18 percentage points higher. Interestingly, the difference between these two groups is relatively stable over different degrees of product complexity. Additionally, we find that shocks to cost overruns and bids are negatively correlated: The estimated correlation for local contractors is -0.142, while for nonlocals it is -0.152. These correlations emphasize the possible drawbacks of promoting too much competition for contracts, as lower bid draws are more likely to produce higher overruns.

Panel C of Table 3 shows that if buyers anticipate that publicity leads to a 10 percent reduction in awarding price, they increase their likelihood of publicizing by 3.7 pp. However, a 10 percent increase in cost overruns reduces the probability of advertising by only 0.4 pp. This asymmetry in the taste parameters for (log) dollars

[52] The procedure to estimate standard errors is discussed in Supplemental Appendix G.2.

[53] It is not surprising that the unobserved heterogeneity term captures a sizable proportion of the price variance, considering that in our data the contract price corresponds to the total awarded value, not distinguishing between the required quantity and unit price. Thus, the unobserved heterogeneity term may capture important quantity differences across contracts. Other unobserved relevant factors are the travel costs between the place of delivery or execution and the bidders' location.

Table 3—Model Estimates: Marginal Effects

			Entry (probit)	Bidding (log normal)	Execution (log normal)	
	Baseline	Change	Marginal effect	Marginal effect	Marginal effect	Marginal effect any ov.
Panel A. Marginal effects						
Service	0.31	1.0	−0.004	0.000	0.014	0.002
Degree of complexity	0.07	0.1	−0.045	−0.007	0.339	0.054
Nonlocal		1.0	0.998	−0.010	0.183	0.023
Nonlocal × complexity		0.1	0.002	−0.004	0.002	0.000
Last month	0.27	1.0	−0.042			
Exp. duration > median	0.50	1.0			0.310	0.037
N^{L}	8.66	1.0	−0.009	−0.002		
N^{NL}	4.45	1.0	−0.165	−0.004		
Panel B. Standard deviation						
Standard deviation				0.099	0.935	
Correlation local				−0.142		
Correlation nonlocal				−0.152		
Standard deviation unobserved het.				0.422		
			Publicity choice (probit)			
Panel C. Marginal effects						
Expected price		0.1	−0.037			
Expected cost overruns		0.1	−0.004			
Expected local winning		0.1	0.017			
Above $25,000		1.0	0.431			

Notes: This table shows model estimates. Panel A shows the marginal effects of different covariates on the main dependent variables related to bidders' actions. Covariates are listed in the first column. Marginal effects are computed at the mean of each covariate, shown in the second column. The third column shows the change in the covariate used to estimate the marginal effect. Columns 4 through 7 show the value of the marginal effects on different dependent variables: the probability of entry, the bid level, the amount of cost overruns, and the probability of any positive overruns. Panel B shows the estimated standard deviation of the normalized log bids, overruns, and the estimated standard deviation of the unobserved cost heterogeneity component. Also, it shows the estimated correlation between normalized log bids and overruns. Panel C displays the marginal effects of four variables on the probability of publicizing the contract solicitation in FBO. The variables are expected log-award price, expected log-cost overruns, expected probability of a local contractor winning, and being above $25,000 in expected award price without publicity. These coefficients are jointly estimated using the SMM .

at the award stage and cost overruns ex post is interesting. One possible explanation is that buyers may underestimate the value of contract execution when they do not deal directly with the contractor during the execution stage.[54]

Buyers have a preference for local vendors. Anticipating a 10 percentage point reduction in the likelihood of a local contractor winning leads buyers to reduce the probability of publicity by 1.7 pp. This coefficient is meaningful in magnitude: The buyer is indifferent between increasing the probability of a local winner by 10 percent and reducing the award price by roughly 4.6 pp. Finally, predicting that the price without advertising will exceed $25,000 increases the likelihood of publicity by 43 pp.

[54] In interviews with procurement officers, we noted that some agencies separate personnel in charge of the contract award and contract management phases of the procurement process.

Model Fit.—Overall, the model replicates the key empirical patterns in the estimation sample. Supplemental Appendix Figure A17 compares the distribution of model-simulated equilibrium outcome variables with actual data. The simulated data resemble actual publicity choices, prices, the number of bidders, and the share of contracts assigned to local contractors. Panels E and F separate cost overruns by products and services. We find that the model slightly overpredicts the jump in publicity at the threshold, which leads to a slight overprediction of the dip in the probability of a local winning the contracts. Also, for services, the model slightly underpredicts the probability of having any (positive) cost overrun but overpredicts the magnitude of cost overruns. This dichotomy suggests that buyers may face frictions when introducing contract modifications ex post that our model does not account for.

Sensitivity.—In Supplemental Appendix H, we explore the sensitivity of our parameter estimates to the estimation moments, following the approach proposed by Andrews, Gentzkow, and Shapiro (2017). Results are consistent with our identification arguments in Section IIIC. Focusing on a few parameters as examples, we show that these are most sensitive to the moments deemed a priori as important for identification.

Recovering Cost Distributions:

Project Costs.—We recover the distribution of project costs by leveraging the methodology introduced by Guerre, Perrigne, and Vuong (2000) and Campo, Perrigne, and Vuong (2003). This method combines the first-order conditions (equation 6)—subject to boundary conditions—with estimated log bids (net of unobserved heterogeneity) to recover the costs that generated these bids. In our setting, the first-order condition depends on the probabilities of different combinations of the number of local and nonlocal bidders. These probabilities are formed from simulations based on the model's participation parameters. Supplemental Appendix Figure A18 shows the probability density function of log costs $\log(\tilde{c}_{jt})$ for both types. Local bidders have slightly higher costs than nonlocals. Supplemental Appendix Figure A19 displays the mean relation between $\tilde{b}_j$ and $\tilde{c}_j$. As expected, markups decrease with higher cost draws.

Entry Costs.—We recover type-specific distributions of entry costs using the equilibrium conditions for optimal entry behavior discussed in Section IIIA. A potential bidder compares the draw from the entry-cost distribution $H^k(\omega)$ with the expected utility of entering—that is, $\varphi^k(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t) = H^k\Big(\bar{\Pi}^k(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)\Big)$. Our estimated project cost distributions $F_c^k(c)$ allow us to form (ex ante) expected utility of participating, $\bar{\Pi}^k(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$, in equation (8), and relate it to observed entry behavior (Athey, Levin, and Seira 2011).

The entry decisions differ substantially across types. On the one hand, nonlocal firms enter with probability 1 for a range of combinations of $(\mathbf{x}_t, \mathbf{z}_t, \mathbf{N}_t)$. Entry choices, combined with the expected utility of entry, yield that nonlocals face (near) zero entry costs for over 80 percent of contracts, and virtually all entry cost draws are smaller than \$1,000. On the other hand, local firms are substantially less likely to enter. Indeed, we find that their entry costs are greater than zero for all their contracts and, with an 80 percent chance, they would not enter for any of the values in the estimated range of expected utilities. This estimated entry-cost asymmetry shapes the composition of actual bidders and, subsequently, the winning bids. Supplemental Appendix Figure A20, panel A shows the probability of entry as a function of the number of nonlocal potential entrants: Locals' modest entry probability shrinks rapidly to zero as competition from nonlocals increases, whereas nonlocals enter with very high probability unless they expect several ($>$ 6) potential rivals. Supplemental Appendix Figure A20, panel B, shows how these entry probabilities translate into the composition of actual competitors (and the winner's identity) as a function of the number of potential nonlocal contractors.

Relationship between Project Costs and Cost Overruns.—Our modeling structure allows project costs $\tilde{c}$ and cost overruns q to be correlated. We achieve this in our estimation by allowing for a correlation between bid shocks and cost overrun shocks, which we can estimate from the inversion of simulated bids and cost overruns. Supplemental Appendix Figure 21 shows mean cost overruns for different levels of project costs, separately for locals and nonlocals, and by product type (goods versus services). Both locals and nonlocals feature a relatively flat relationship between overruns and project costs for goods, but the relationship is negative for services, meaning that lower cost draws are associated with higher cost overruns. Relative to locals, nonlocals exhibit a more pronounced negative relationship between cost overruns and project costs and higher levels of cost overruns for any given cost. In terms of magnitudes, a 30 percentage point lower cost draw (equivalent to the twenty-fifth to fifth percentile gap) is partially offset by a 7 percentage point increase in cost overruns for service contracts by nonlocals. We return to the implications of these patterns in Section IVA.

IV. Counterfactual Analysis

We use our model to evaluate the implications of counterfactual scenarios. In the first counterfactual exercise, we gauge the effects of increasing competition through publicity on contract outcomes. This exercise extends and complements the reduced form estimates obtained in Section II and allows us to assess quantitatively the mechanisms that drive the effects we observe in the data. The second and third counterfactuals evaluate the implications of modifying publicity rules. In the second counterfactual, we analyze the consequences of removing publicity altogether, allowing every buyer to decide whether or not to advertise the solicitation. The third exercise studies the effects of an alternative regulatory design that makes publicity requirements a function of the underlying contract complexity.

A. *Effects of Increasing Competition through Publicity*

Our model estimates allow us to replicate and extend the results from Section II. Using the model, we can evaluate contract outcomes with and without publicity throughout the range of awards in our sample, not only for contracts around the threshold. Indeed, relative to the reduced form local average treatment effects presented before, the estimated model allows us to measure publicity effects over the complete sample—that is, estimate the average treatment effect of publicity.

Figure 7 shows a series of outcomes as a function of the expected contract price. For the current policy (red line), results are broadly in line with our reduced-form analysis: The number of bidders jumps at the threshold under the current policy (panel A), allowing nonlocals to participate and win at higher rates (panel B), slightly reducing award prices (panel C), and increasing cost overruns (panel D). Our model allows us to benchmark the observed regime with three alternative scenarios: one with no publicized contracts (no publicity), one with all publicized contracts (full publicity), and one where officers are free to choose throughout the expected price range (no-threshold regulation). Intuitively, the no-threshold regulation is equivalent to the current policy below the threshold and extends this regime to awards expected to exceed it, eliminating the discontinuities. On the other hand, full publicity and no publicity represent bounds on the extent to which the regulator can leverage publicity to affect outcomes. Full publicity maximizes the number of bidders and minimizes the share of local winners, leading to the lowest prices and the highest overruns. The opposite is true for no publicity.

To assess whether ex ante price effects or ex post cost overruns effects dominate, we consider the final price $P^{D}_{F,t}$:

$$P^{D}_{F,t} = P^{D}_{I,t}\left(1 + Q^{D}_{t}\right), \tag{12}$$

where $D_t \in \{0,1\}$ denotes contract t's publicity status, $P^{D}_{I,t}$ is the award price, and Q^{D}_{t} is the share of cost overruns ex post.

Figure 8 compares contract award prices, cost overruns, and final prices with and without publicity for different levels of complexity. The gray line at zero marks the no-publicity benchmark. The orange dashed line shows the effects of publicity on award price, and the green dashed line shows the effect on cost overruns. Both effects are more pronounced for complex contracts. This is intuitive for cost overruns but less obvious for award prices and is explained by two factors. First, auctions that require complex contracts have a higher variance in bid functions, which increases the support of possible price reductions from additional bidders. Second, auctions for complex contracts have lower participation, meaning that the effect of an extra bidder is higher than when there are fewer competitors. The blue solid line shows the combined effect of ex ante price reductions and ex post price increases. Publicity reduces total prices for contracts with low levels of complexity but increases them for more complex acquisitions.

These findings provide quantitative evidence consistent with the idea introduced by seminal papers on incomplete contracting: There exists a degree of transaction complexity beyond which promoting competition may backfire. When there is a high number of possible contingencies during the execution stage, ensuring adequate performance ex post may be more important than reducing prices ex ante. (Williamson 1976; Bajari and Tadelis 2001; Bajari et al. 2014; Bolotnyy and Vasserman 2023).

Breaking Down the Effects on Contract Performance.—In Section IIIE, we show that our estimates imply a negative relationship between contract completion costs

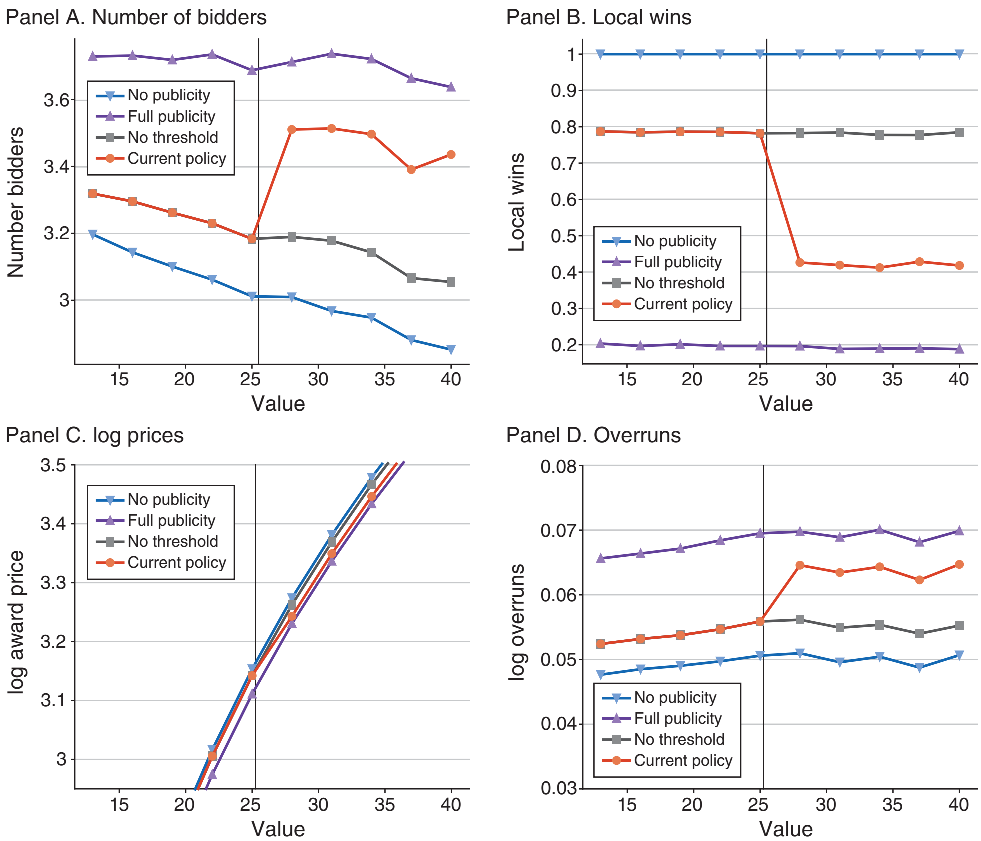

Figure 7. Policy Evaluation Using Model Estimates

Notes: This figure shows different outcome variables around the threshold for different counterfactual policies. Panel A shows the number of bidders, panel B the probability of awarding to a local vendor, panel C the log award price, and panel D log overruns. The x-axis in every graph is the expected award value of the contract without publicity. In every graph, we compare the current policy design (red line) with counterfactual regulations mandating full and zero publicity. We also include a counterfactual policy of no threshold, implying that the buyers may freely choose whether to publicize contracts throughout the expected award range.

and cost overruns and that the strength of this relationship is type-dependent. This fact suggests that the effects of competition on performance arise due to a combination of within-type and cross-type effects. Within-type effects occur since winners in more competitive auctions tend to draw costs from the lower part of the cost distribution, and the estimated negative correlation implies that this will lead to higher cost overruns. Cross-type effects instead come from the increased participation of nonlocal firms, which tend to draw worse performance shocks.

We leverage the model estimates to disentangle these effects quantitatively. Two aspects of this exercise are worth highlighting. First, this investigation is related to the reduced form analysis presented in Section IIE, where we show that the effects of publicity on contract performance were primarily driven by cross-firm rather than within-firm variation. Second, the results may be informative about the plausibility of some of our modeling choices. Finding a large role for within-type effects could indicate model misspecification reflecting, for example, that our parametrization of

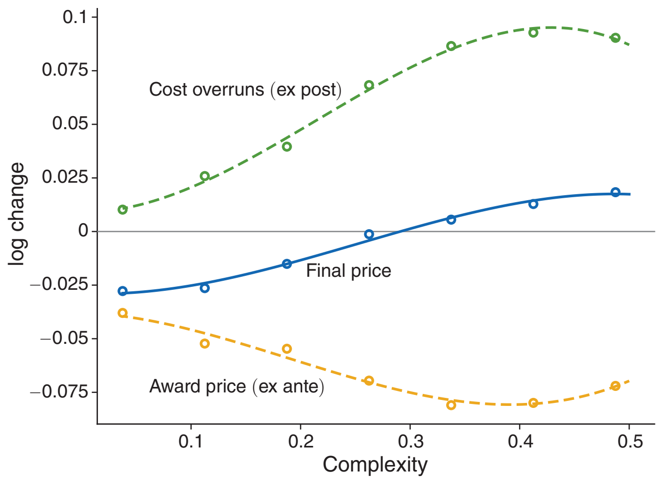

FIGURE 8. EFFECTS OF PUBLICITY ON EX ANTE, EX POST, AND FINAL PRICES

Notes: This figure shows the effect of publicity on (log) award prices, (log) cost overruns, and (log) final prices, as a function of product complexity. Note that from the final price definition (12), we have that $log(P_{F,t}^{D}) = log(P_{I,t}^{D}) + log(1 + Q_{t}^{D})$, so the effect on log final price is the addition of the effect on log award price and the effect on log overruns. The figure presents the effects of publicity relative to a benchmark of no publicity, represented in the horizontal line at zero. Circles exhibit the mean effect by complexity bin. Each line corresponds to a flexible polynomial fit. The degree of complexity is defined as the log of the product category's average overruns for contracts below \$20,000.

heterogeneity into two groups is too restrictive or that strategic considerations that we have shut down by assumption may be relevant in practice. Instead, finding that most of the influence of competition on performance is due to cross-type effects would be reassuringly consistent with the reduced form evidence and our simplifying assumptions regarding the source of performance differences.

We consider the following decomposition:

$$(13)\quad \underbrace{log(1 + Q_{t}^{1}) - \log(1 + Q_{t}^{0})}_{\text{Total effect}} = \underbrace{log(1 + Q_{t}^{1}) - log(1 + Q_{t}^{1L})}_{\text{Cross-type effect}} + \underbrace{log(1 + Q_{t}^{1L}) - log(1 + Q_{t}^{0})}_{\text{Within-type effect}},$$

where Q_{t}^{0} corresponds to the level of cost overruns if contract t is not publicized, Q_{t}^{1} is the level of cost overruns if the contract is publicized, and Q_{t}^{1L} is the level of cost overruns if the contract is publicized and we force the firm that executes the contract to draw overrun shocks from the distribution of locals, keeping all other factors constant. The cross-type effect captures changes in overruns due to changes in the contractor's group affiliation, while the within-type effect captures changes in overruns within the group of locals. Note that when a publicized contract is won by a local, the cross-type effect is mechanically zero ($Q_{t}^{1} = Q_{t}^{1L}$), and the full effect is due to changes within a type.

We implement this decomposition and find that 77 percent of the increase in overruns is explained by cross-type effects. This implies that if we shut down cross-type effects, enhanced competition would induce less than one-fourth of the increase in overruns we observe in the data. Moreover, we find that the relative contribution of these effects is heterogeneous, with cross-type effects accounting for 89 percent and 70 percent of the total effects for goods and services, respectively. Figure 9 shows the decomposition of the impact on overruns for different degrees of contract complexity. Even though the contribution of cross-type effects slightly decreases with complexity, it remains the main factor explaining the increase in overruns. We interpret these results as broadly consistent with our reduced form evidence and supportive of our key modeling assumptions regarding performance effects.

B. *The Value of Delegating Competition Promotion to the Buyer*

We now assess the implications of allowing the buyer to choose whether to openly publicize a contract, as opposed to mandating rules that constrain the buyer's discretion. This trade-off pertains to the more general problem of the delegation of authority within organizations (Aghion and Tirole 1997) and has been frequently analyzed in the context of public procurement.[55]

Conceptually, the publicity requirement acts as a discontinuous jump in the cost of not publicizing. Below the threshold, buyers choose whether to advertise the contract solicitation with full discretion; above the threshold, regulation forces them to publicize more often than desired. Using the estimated model parameters, we simulate buyers' hypothetical decisions in a no threshold situation, with full discretion over publicity decisions for the whole range of contracts in our sample. Like before, we benchmark this counterfactual against the current policy design, a scenario without publicity for any contract, and a full publicity regulation where all contracts are publicized regardless of their size.

Figure 10 displays the results. Panel A shows the fraction of publicized contracts as a function of award values. Panel B shows changes in log final prices relative to no publicity (the benchmark gray line at zero) and as a function of product complexity. We find that the effects of discretion are ambiguous compared to no publicity or full publicity. On the one hand, relative to no publicity, discretion (no threshold) leads to lower final contracting costs for contracts with a degree of complexity below 0.34—as a reference, 0.34 corresponds to the complexity of contracts for "housekeeping–laundry/dry cleaning." A no-publicity rule outperforms discretion for more complex contracts because the buyer underweights overruns relative to the award price. On the other hand, a full publicity rule delivers larger cost savings than discretion for contracts with sufficiently low levels of complexity. Finally, for an intermediate range of complexity values, discretion yields lower contract costs than the two extreme scenarios. These results suggest ample space for improvements to the current regulation, as the value of discretion heavily depends on the transaction's underlying degree of complexity. Supplemental Appendix I discusses the role of buyer preference parameters in shaping contract outcomes.

[55] See, e.g., Kelman (1990); Coviello et al. (2018); Carril (2022); Szucs (2024); Bandiera et al. (2021); Bosio et al. (2020); Decarolis et al. (2025).

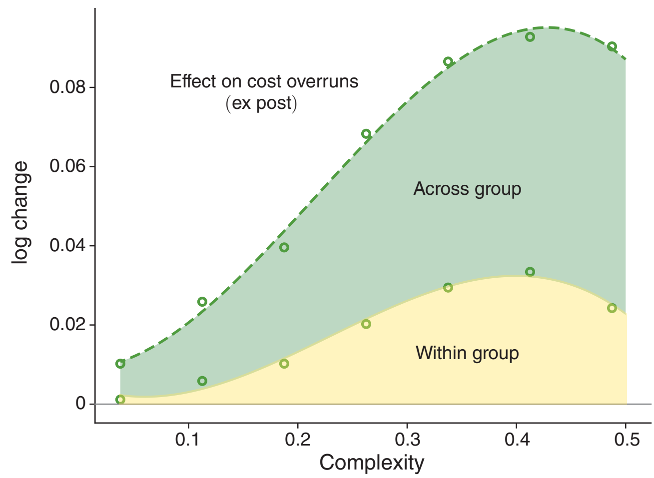

FIGURE 9. DECOMPOSITION OF THE EFFECTS OF PUBLICITY ON COST OVERRUNS

Notes: This figure shows the effect of publicity on (log) cost overruns, as a function of product complexity. The total effect of publicity compares the level of log overruns with publicity relative to a benchmark of no publicity, represented in the horizontal line at zero. The within-type and across-type decomposition follows (13). Circles exhibit the mean effect by complexity bin. Each line corresponds to a flexible polynomial fit. The degree of complexity is defined as the log of the product category's average overruns for contracts below $20,000.

C. *Complexity-Based Publicity Requirements*

We now consider varying the regulation design depending on the level of contract complexity. The proposed exercise contemplates identifying the cost-minimizing level of publicity requirements for each degree of complexity. By publicity requirements, we mean the share of contracts that buyers must publicize. Even though other, more sophisticated regulatory tools may enhance procurement efficiency, the design of real-world procurement rules faces a constraint on the level of intricacy. Therefore, we choose to analyze counterfactual complexity-based requirements, considering that they would represent relatively minor modifications to the current environment and thus could be realistically implemented.[56] We show that even this minor regulatory change may result in substantial procurement cost savings.

We proceed in three steps. First, we simulate contract outcomes under different levels of product-specific publicity requirements that buyers are mandated to meet. Second, we estimate the final price under each of these requirements. Finally, we identify the publicity requirement that yields the lowest final price at each complexity level.

Figure 11 summarizes this procedure. Panel A illustrates the publicity requirement that minimizes the final price at different complexity levels. Panel B shows

[56]The current version of FAR part 5 ("Publicizing Contract Actions") allows buyers to apply for exemptions if they prefer not to publicize a contract. The proposed policy design could be implemented by simply varying the set of exemptions that different product categories are allowed to invoke. For example, if the contract solicitation involves a well-defined product for which the product category's cost-minimizing level of publicity requirement is 100 percent, then there would be no exemption to be invoked. Conversely, if the solicitation requires a relatively complex product, the buyer could have more (or total) discretion to file exemptions.

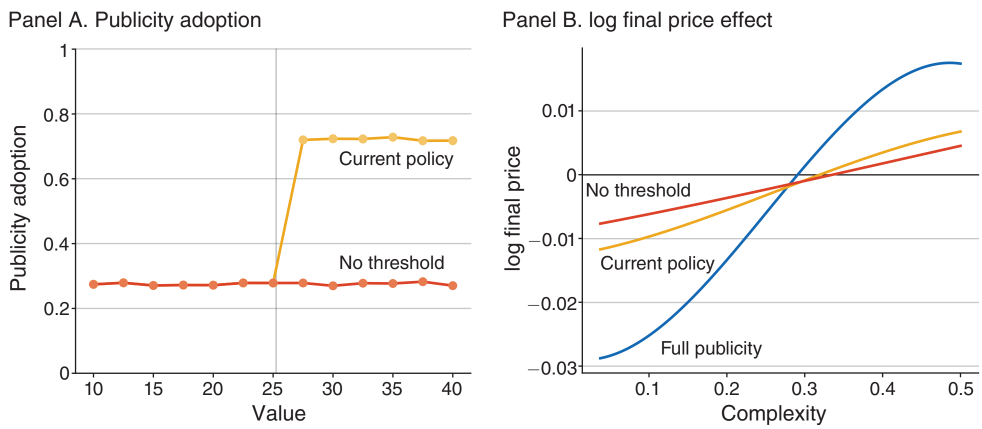

Figure 10. Delegation of Publicity Decision to the Buyer

Notes: Panel A shows the share of publicized contracts as a function of expected award value in the absence of publicity. The orange line displays the current policy with a regulation threshold at \$25,000. The red line describes a counterfactual with no regulation threshold, meaning that each buyer can freely decide whether to publicize contracts throughout the expected award range. Panel B shows changes in log final prices as a function of product complexity, relative to a benchmark of no publicity, and for different counterfactual policies. The blue line is a counterfactual where all contracts are publicized, the orange line represents the current policy (with a threshold at \$25,000), and the red line shows the no-regulation threshold counterfactual. Each line corresponds to a flexible polynomial fit. The degree of complexity is defined as the log of the product category's average overruns for contracts below \$20,000.

the effect of complexity-based publicity requirements on final prices (brown-dashed line), relative to a benchmark of no publicity, and compares it with the effects of full discretion (no threshold) and full publicity. The cost-minimizing publicity requirements are close to 100 percent for low-complexity products, and the optimal share of publicity requirements falls gradually as complexity increases. Panel B shows this yields final price reductions relative to the other counterfactual scenarios. When the product purchased is simple, the tailored publicity requirements remove buyers' discretion entirely to leverage the benefits of enhanced competition. However, it provides more discretion when contracts are more complex to attenuate the negative consequences of contract implementation ex post.

Relative to a baseline of no publicity for any contract, the current policy design with publicity requirements at \$25,000 reduces, on average, the final price by 0.94 percent. The counterfactual policy of tailoring the publicity requirements to the degree of complexity outperforms all the other considered regulations, reducing average prices by 2.39 percent. The 1.45 percentage point difference in cost savings between the current regime and the complexity-based design corresponds to \$103.5 million per year.[57]

[57] This amount is calculated extrapolating to all competitively awarded definite contracts from the DOD in 2019 with values below the simplified acquisition threshold (i.e., \$250,000).

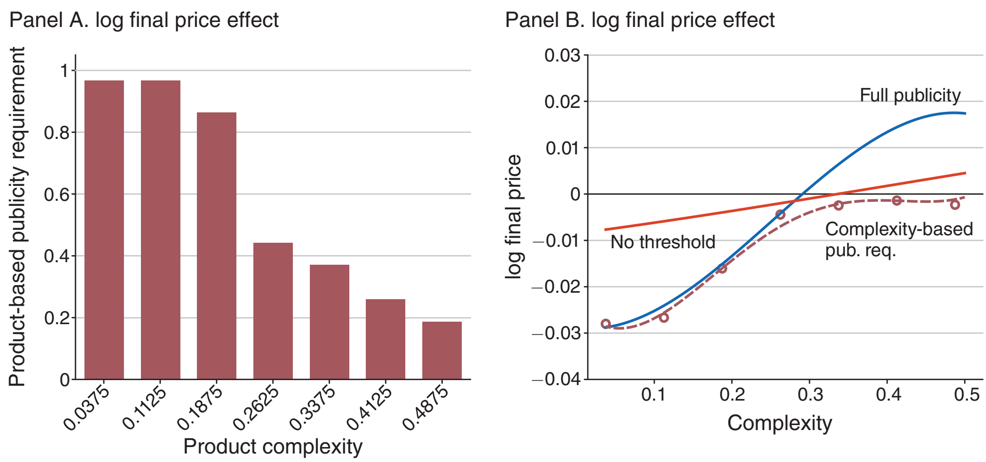

Figure 11. Counterfactual Analysis III: Complexity-Based Requirements

Notes: Panel A shows the level of publicity requirement that yields the minimum final prices for different levels of complexity. Panel B shows the effects on final cost relative to the no-publicity benchmark at each degree of complexity. The blue line is a counterfactual where all contracts are publicized, the red line shows the no regulation threshold counterfactual, and the brown dashed line corresponds to the price effect under the cost-minimizing publicity requirement. Each line corresponds to a flexible polynomial fit. The brown dots correspond to the mean at different complexity bins. The degree of complexity is defined as the log of the product category's average overruns for contracts below $20,000.

D. *Discussion*

While the main regulatory counterfactual we explore is slightly more sophisticated than the current uniform policy, it is still remarkably simple and keeps most institutional features unchanged. We do this purposely to emphasize the nontrivial cost savings of even a minor, easy-to-implement change. While estimating the impact of more involved policy reforms exceeds the scope of this study, two alternative policy designs are worth mentioning.

The first is to change the auction design to account for adverse selection. Various approaches have been proposed in the literature. One recent example is Lopomo, Persico, and Villa (2023), who show that the optimal mechanism in the presence of noncontractible quality concerns features both a price ceiling and, more surprisingly, a price floor. The mechanism, which the authors call a lowball lottery auction, uses the price floor to reduce the advantage of cheap suppliers, limiting the extent to which they can bid aggressively. In our setting, a price floor in more complex product categories may reduce total costs by reducing the likelihood of a low-performance winner.

A second approach would be explicitly including past performance measures in the awarding mechanism.[58] Our results suggest that considering past performance may counteract some of the negative performance effects we document since they may help screen out low-cost, low-performance contractors. Still, the full effects of this approach may be more subtle since its impact on bidding strategies will depend on the interplay

[58] Past performance is recorded and used in DOD procurement to evaluate proposals in awards exceeding the simplified acquisition threshold (currently, $250,000). The evaluations for these awards are recorded in the CPARS.

between costs and performance and their heterogeneity (Andreyanov et al. 2023). Finally, using past performance may act as a reputational incentive in settings where moral hazard is a dominating force (Calzolari and Spagnolo 2009; Malcomson 2013; Spagnolo 2012). For example, Coviello, Guglielmo, and Spagnolo (2018) find evidence that the use of discretion by public buyers in Italy increases the probability of repeated winners and that this credible commitment to possible repeated contracting disciplines suppliers into having high performance.

V. Conclusion

This paper studies the relationship between competition and procurement contract outcomes. Even though procurement represents a significant share of the economy, there is limited evidence on the implications of policies oriented to expand competition, considering not only the effects on award prices but also on post-award contract execution. Using data and policy variation from DOD procurement, we provide extensive evidence on the effects of increasing competition through publicity.

Our identification strategy leverages a regulation that generates quasiexperimental variation in the extent to which contract opportunities are broadly advertised to potential suppliers. We find that contract publicity increases the number of participating bidders and that this added competitive pressure results in lower acquisition prices. However, broader advertisement leads to a different pool of vendors, who perform worse ex post. Our analysis shows that the degree of contract complexity determines the extent of this trade-off. Promoting competition reduces total contract costs only for simple transactions, as relatively complex ones are more exposed to cost overruns and delays in the execution stage.

Motivated by this evidence, we develop and estimate an equilibrium model of competitive bidding for procurement contracts. The model allows us to estimate the underlying firms' characteristics that shape adverse selection in this market and buyers' objectives when promoting competition through advertising. We use our estimates to evaluate relevant counterfactual policies. Our results show that delegating competition promotion to the buyer has ambiguous effects. For an intermediate range of complexity values, discretion yields lower contract costs than simple rules of zero or mandatory publicity, which suggests ample space for improvements to the current threshold-based publicity rules. Thus, we use our model to engineer improvements to the current policy design by introducing publicity requirements tailored to the degree of complexity of the purchase. We find that departing from a uniform regulation would significantly reduce procurement costs.

While our empirical analysis focuses on DOD procurement, our results have wider applicability. The highly specific nature of defense acquisitions is readily apparent for major contracts involving fighter jets or weapons systems. Yet, our sample of contracts below $40,000 is composed of products and services similar to those procured by civilian agencies and the private sector, including software and computers, medical supplies, furniture, maintenance services, and utilities and housekeeping.[59] Moreover, while small awards represent a modest share of all federal procurement

[59] See Supplemental Appendix Table B.2.

spending, they do reflect the typical contracting activity of the federal government.[60] Therefore, we expect the tension between price competition and contract performance that we study in this paper will be relevant in many other procurement settings.

REFERENCES

Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy* 105 (1): 1–29.

Albano, Gian Luigi, Milo Bianchi, and Giancarlo Spagnolo. 2006. "Bid Average Methods in Procurement." *Rivista di Politica Economica* 96 (1): 41–62.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *Quarterly Journal of Economics* 132 (4): 1553–92.

Andreyanov, Pasha, Francesco Decarolis, Riccardo Pacini, and Giancarlo Spagnolo. 2023. "Past Performance and Procurement Outcomes." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4929595.

Asker, John, and Estelle Cantillon. 2008. "Properties of Scoring Auctions." *RAND Journal of Economics* 39 (1): 69–85.

Asker, John, and Estelle Cantillon. 2010. "Procurement when Price and Quality Matter." *RAND Journal of Economics* 41 (1): 1–34.

Athey, Susan. 2001. "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information." *Econometrica* 69 (4): 861–89.

Athey, Susan, Dominic Coey, and Jonathan Levin. 2013. "Set-Asides and Subsidies in Auctions." *American Economic Journal: Microeconomics* 5 (1): 1–27.

Athey, Susan, Jonathan Levin, and Enrique Seira. 2011. "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions." *Quarterly Journal of Economics* 126 (1): 207–57.

Bajari, Patrick. 2001. "Comparing Competition and Collusion: A Numerical Approach." *Economic Theory* 18 (1): 187–205.

Bajari, Patrick, Stephanie Houghton, and Steve Tadelis. 2014. "Bidding for Incomplete Contracts: An Empirical Analysis." *American Economic Review* 104 (4): 1288–319.

Bajari, Patrick, Robert McMillan, and Steven Tadelis. 2009. "Auctions versus Negotiations in Procurement: An Empirical Analysis." *Journal of Law, Economics, and Organization* 25 (2): 372–99.

Bajari, Patrick, and Steven Tadelis. 2001. "Incentives versus Transaction Costs: A Theory of Procurement Contracts." *RAND Journal of Economics* 32 (3): 387–407.

Bandiera, Oriana, Michael Carlos Best, Adnan Qadir Khan, and Andrea Prat. 2021. "The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats." *Quarterly Journal of Economics* 136 (4): 2195–242.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti. 2009. "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment." *American Economic Review* 99 (4): 1278–308.

Barreca, Alan I., Melanie Guldi, Jason M. Lindo, and Glen R. Waddell. 2011. "Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification." *Quarterly Journal of Economics* 126 (4): 2117–23.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi. 2023. "Individuals and Organizations as Sources of State Effectiveness." *American Economic Review* 113 (8): 2121–67.

Bhattacharya, Vivek, James W. Roberts, and Andrew Sweeting. 2014. "Regulating Bidder Participation in Auctions." *RAND Journal of Economics* 45 (4): 675–704.

Bolotnyy, Valentin, and Shoshana Vasserman. 2023. "Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement." *Econometrica* 91 (4): 1205–59.

Bosio, Erica, Simeon Djankov, Edward Glaeser, and Andrei Shleifer. 2020. "Public Procurement in Law and Practice." *American Economic Review* 112 (4): 1091–117.

Brendstrup, Bjarne, and Harry J. Paarsch. 2003. "Nonparametric Estimation of Dutch and First-Price, Sealed-Bid Auction Models with Asymmetric Bidders." Unpublished.

Bresnahan, Timothy F., and Peter C. Reiss. 1991. "Entry and Competition in Concentrated Markets." *Journal of Political Economy* 99 (5): 977–1009.

Bulow, Jeremy, and Paul Klemperer. 1996. "Auctions versus Negotiations." *American Economic Review* 86 (1): 180–94.

[60] In fiscal year 2017, 71 percent of the contracts awarded by non-DOD federal agencies were valued at $40,000 or less.

Burguet, Roberto, Juan José Ganuza, and Esther Hauk. 2012. "Limited Liability and Mechanism Design in Procurement." *Games and Economic Behavior* 76 (1): 15–25.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–326.

Calzolari, Giacomo, and Giancarlo Spagnolo. 2009. "Relational Contracts and Competitive Screening." CEPR Discussion Paper 7434.

Campo, Sandra, Isabelle Perrigne, and Quang Vuong. 2003. "Asymmetry in First-Price Auctions with Affiliated Private Values." *Journal of Applied Econometrics* 18 (2): 179–207.

Carril, Rodrigo. 2022. "Rules versus Discretion in Public Procurement." Unpublished.

Carril, Rodrigo, Andres Gonzalez-Lira, and Michael S. Walker. 2026. *Data and Code for: "Competition under Incomplete Contracts and the Design of Procurement Policies."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. Https://doi.org/10.3886/E231127V1.

Cattaneo, Matias D., and Rocío Titiunik. 2022. "Regression Discontinuity Designs." *Annual Review of Economics* 14: 821–51.

Chillemi, Ottorino and Claudio Mezzetti. 2014. "Optimal Procurement Mechanisms: Bidding on Price and Damages for Breach." *Economic Theory* 55 (2): 335–55.

Compiani, Giovanni, Philip Haile, and Marcelo Sant'Anna. 2020. "Common Values, Unobserved Heterogeneity, and Endogenous Entry in US Offshore Oil Lease Auctions." *Journal of Political Economy* 128 (10): 3872–912.

Conley, Timothy G., and Francesco Decarolis. 2016. "Detecting Bidders Groups in Collusive Auctions." *American Economic Journal: Microeconomics* 8 (2): 1–38.

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

Coviello, Decio, Andrea Guglielmo, and Giancarlo Spagnolo. 2018. "The Effect of Discretion on Procurement Performance." *Management Science* 64 (2): 715–38.

Coviello, Decio, and Mario Mariniello. 2014. "Publicity Requirements in Public Procurement: Evidence from a Regression Discontinuity Design." *Journal of Public Economics* 109: 76–100.

Davezies, Laurent, and Thomas Le Barbanchon. 2017. "Regression Discontinuity Design with Continuous Measurement Error in the Running Variable." *Journal of Econometrics* 200 (2): 260–81.

Decarolis, Francesco. 2014. "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions." *American Economic Journal: Applied Economics* 6 (1): 108–32.

Decarolis, Francesco. 2018. "Comparing Public Procurement Auctions." *International Economic Review* 59 (2): 391–419.

Decarolis, Francesco, Raymond Fisman, Paolo Pinotti, and Silvia Vannutelli. 2025. "Rules, Discretion, and Corruption in Procurement: Evidence from Italian Government Contracting." *Journal of Political Economy Microeconomics* 3 (2): 213–54.

Decarolis, Francesco, Leonardo M. Giuffrida, Elisabetta Iossa, Vincenzo Mollisi, and Giancarlo Spagnolo. 2020. "Bureaucratic Competence and Procurement Outcomes." *Journal of Law, Economics, and Organization* 36 (3): 537–97.

Dini, Federico, Riccardo Pacini, and Tommaso Valletti. 2006. "Scoring Rules." In *Handbook of Procurement*, edited by Nicola Dimitri, Gustavo Piga, Giancarlo Spagnolo, 293–321. Cambridge University Press.

Espín-Sánchez, José-Antonio, Álvaro Parra, and Yuzhou Wang. 2021. "Equilibrium Uniqueness in Entry Games with Private Information." Cowles Foundation Discussion Paper 2126R.

Gerard, François, Miikka Rokkanen, and Christoph Rothe. 2020. "Bounds on Treatment Effects in Regression Discontinuity Designs with a Manipulated Running Variable." *Quantitative Economics* 11 (3): 839–70.

Gourieroux, C., A. Monfort, and E. Renault. 1993. "Indirect Inference." *Journal of Applied Econometrics* 8 (S1): S85–S118.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica* 68 (3): 525–74.

Haile, Philip A., and Yuichi Kitamura. 2019. "Unobserved Heterogeneity in Auctions." *Econometrics Journal* 22 (1): C1–C19.

Hong, Han, and Matthew Shum. 2002. "Increasing Competition and the Winner's Curse: Evidence from Procurement." *Review of Economic Studies* 69 (4): 871–98.

Jae Eun, Dong. 2018. "Reducing Ex-Post Adjustment by Bid Screening Ex-Ante: Evidence from Procurement Auctions for Paving Work." Unpublished.

Kang, Karam, and Robert A. Miller. 2021. "Winning by Default: Why Is There so Little Competition in Government Procurement?" *Review of Economic Studies* 89 (3): 1495–556.

Kelman, Steven. 1990. *Procurement and Public Management: The Fear of Discretion and the Quality of Government Performance*. AEI Press.

Krasnokutskaya, Elena. 2011. "Identification and Estimation of Auction Models with Unobserved Heterogeneity." *Review of Economic Studies* 78 (1): 293–327.

Krasnokutskaya, Elena, and Katja Seim. 2011. "Bid Preference Programs and Participation in Highway Procurement Auctions." *American Economic Review* 101 (6): 2653–86.

Krishna, Vijay. 2003. *Auction Theory*. Elsevier.

Lebrun, Bernard. 1999. "First Price Auctions in the Asymmetric N Bidder Case." *International Economic Review* 40 (1): 125–42.

Lee, David S., and Thomas Lemieux. 2010. "Regression Discontinuity Design in Economics." *Journal of Economic Literature* 48 (2): 281–355.

Levin, Dan, and James L. Smith. 1994. "Equilibrium in Auctions with Entry." *American Economic Review* 84 (3): 585–99.

Lewis, Gregory, and Patrick Bajari. 2011. "Procurement Contracting with Time Incentives Theory and Evidence." *Quarterly Journal of Economics* 126 (3): 1173–211.

Li, Tong, and Xiaoyong Zheng. 2009. "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions." *Review of Economic Studies* 76 (4): 1397–429.

Li, Tong, and Xiaoyong 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, Jeffrey B., and Neale Mahoney. 2017. "Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement." *American Economic Review* 107 (11): 3510–49.

Lopomo, Giuseppe, Nicola Persico, and Alessandro T. Villa. 2023. "Optimal Procurement with Quality Concerns." *American Economic Review* 113 (6): 1505–29.

McFadden, Daniel. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." *Econometrica* 57 (5): 995–1026.

MacKay, Alexander. 2022. "Contract Duration and the Costs of Market Transactions." *American Economic Journal: Microeconomics* 14 (3): 164–212.

Macleod, W. Bentley, and James M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57 (2): 447–80.

Malcomson, James M. 2012. "Relational Incentive Contracts." In *Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 1014–65. Princeton University Press.

Malcomson, James M. 2013. "Relational Incentive Contracts." In *Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 1014–65. Princeton University Press.

Manelli, Alejandro M., and Daniel R. Vincent. 1995. "Optimal Procurement Mechanisms." *Econometrica* 63 (3): 591–620.

Marmer, Vadim, Artyom Shneyerov, and Pai Xu. 2013. "What Model for Entry in First-Price Auctions? A Nonparametric Approach." *Journal of Econometrics* 176 (1): 46–58.

Maskin, Eric, and John Riley. 2003a. "Asymmetric Auctions." *Review of Economic Studies* 67 (3): 413–38.

Maskin, Eric, and John Riley. 2003b. "Equilibrium in Sealed High Bid Auctions." *Review of Economic Studies* 67 (3): 439–54.

Pakes, Ariel, and David Pollard. 1989. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica* 57 (5): 1027–57.

Pei, Zhuan, and Yi Shen. 2017. "The Devil Is in the Tails: Regression Discontinuity Design with Measurement Error in the Assignment Variable." *Advances in Econometrics* 38: 455–502.

Porter, Robert H., and J. Douglas Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *Journal of Political Economy* 101 (3): 518–38.

Roberts, James W. 2013. "Unobserved Heterogeneity and Reserve Prices in Auctions." *RAND Journal of Economics* 44 (4): 712–32.

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

Samuelson, William F. 1985. "Competitive Bidding with Entry Costs." *Economics Letters* 17 (1–2): 53–57.

Spagnolo, Giancarlo. 2012. "Reputation, Competition, and Entry in Procurement." *International Journal of Industrial Organization* 30 (3): 291–96.

Spulber, Daniel F. 1990. "Auctions and Contract Enforcement." *Journal of Law, Economics, and Organization* 6 (2): 325–44.

Storn, Rainer, and Kenneth Price. 1997. "Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces." *Journal of Global Optimization* 11 (4): 341–59.

Sweeting, Andrew, and Vivek Bhattacharya. 2015. "Selective Entry and Auction Design." *International Journal of Industrial Organization* 43: 189–207.

Szucs, Ferenc. 2024. "Discretion and Favoritism in Public Procurement." *Journal of the European Economic Association* 22 (1): 117–60.

US Bureau of the Fiscal Service. 2006–2019. *US Federal Procurement Contract Awards*. Inter-university Consortium for Political and Social Research. https://doi.org/10.3886/E237344V1.

US Census Bureau. 2010a. *2010 ZCTA to CBSA Relationship File* (*zcta_cbsa_rel_10.txt*). https://www2.census.gov/geo/docs/maps-data/data/rel/zcta_cbsa_rel_10.txt (accessed May 26, 2018).

US Census Bureau. 2010b. *2010 ZCTA to County Relationship File* (*zcta_county_rel_10.txt*). https://www2.census.gov/geo/docs/maps- data/data/rel/zcta_county_rel_10.txt (accessed May 25, 2018).

US Census Bureau. 2010c. *Gazetteer Files: 2010 ZIP Code Tabulation Areas* (*ZCTA*). https://www.census.gov/geographies/reference-files/2010/geo/gazetter- file.html (accessed July 7, 2018).

US Census Bureau. 2016. *Census Bureau Region and Division Codes and FIPS Codes for States and Counties* (*stategeocodes-v2016.xls*). https://www2.census.gov/programssurveys/popest/geographies/2016/state-geocodes- v2016.xls (accessed October 9, 2018).

US General Services Administration. 2018. *FPDS-NG Valid Contracting Offices* (*FPDSNG_Contracting_Offices.xls*). https://www.fpds.gov/downloads/top_requests/FPDSNG_Contracting_Offices.xls (accessed May 17, 2018).

US General Services Administration. 2020a. *FPDS-NG Valid Contracting Offices* (*FPDSNG_Contracting_Offices.xls*). https://www.fpds.gov/ (accessed 2020).

US General Services Administration. 2020b. *Product and Service Codes* (*PSC*) *Manual*. https://www.acquisition.gov/psc-manual (accessed 2020).

US General Services Administration. 2007–2020. *US Federal Procurement Solicitation Notices from fedbizopps.gov, 2007–2020*. Inter-university Consortium for Political and Social Research.

Williamson, Oliver E. 1976. "Transaction-Cost Economics: The Governance of Contractual Relations." *Journal of Law and Economics* 22 (2): 233–61.