Title
Essays on Markets and Institutions

Author
Gonzalez-Lira, Andres

Publication Date
2021

Peer reviewed|Thesis/dissertation

Essays on Markets and Institutions

by

Andres Gonzalez Lira

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Steve Tadelis, Chair
Professor Benjamin Handel
Professor Kei Kawai
Professor Reed Walker

Spring 2021

Essays on Markets and Institutions

Abstract

Essays on Markets and Institutions

by

Andres Gonzalez Lira

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Steve Tadelis, Chair

This dissertation studies how agents and firms interplay with market institutions. In the first two chapters, co-authored with Rodrigo Carril and Michael Walker, we study the implications of policies oriented to intensify competition for procurement contracts. Conceptually, opening contracts up to bids by more participants leads to lower acquisition costs. However, expanding the set of bidders hinders buyers' control over the quality of prospective contractors, potentially exacerbating adverse selection on non-contractible quality dimensions. We study this trade-off in the context of procurement by the U.S. Department of Defense. Our empirical strategy leverages regulation that mandates agencies to publicize contract opportunities whose value is expected to exceed a certain threshold. We find that advertising contract solicitations increases competition and leads to a different pool of selected vendors who, on average, offer lower prices. However, it also worsens post-award performance, resulting in more cost overruns and delays. This negative effect on post-award performance is driven by goods and services that are relatively complex, highlighting the role of contract incompleteness. To further study the scope of this tension, we develop and estimate a model in which the buyer chooses the extent of competition, and the invited sellers decide on auction participation and bidding. We estimate sellers' cost and ex-post quality distributions, as well as buyers' preference parameters over contract outcomes. Simulating equilibrium conditions under counterfactual settings, we benchmark the current regulation design with complexity-tailored publicity requirements, and find that adjustments to publicity requirements could provide savings of 2 percent of spending, or $104 million annually.

One of the main roles of the government involves enacting and enforcing regulation aimed at curbing the undesired behavior of agents. In the third chapter, co-authored with Mushfiq Mobarak, we study the consequences of deploying enforcement activities over illegal activities in local markets in Chile. The paper grapples with a key real-world feature that is that regulated agents adapt to circumvent enforcement. We present and test a model of enforcement with learning and adaptation by auditing vendors selling illegal fish in Chile in a randomized controlled trial and tracking them daily using mystery shoppers. Leveraging experimental variation on the frequency and predictability of enforcement, we can test the

model's predictions and find that conducting audits on a predictable schedule and (counter-intuitively) at high frequency is less effective, as agents learn to take advantage of loopholes. A consumer information campaign proves to be almost as cost-effective and curbing illegal sales and obviates the need for complex monitoring and policing. The Chilean government subsequently chooses to scale up this campaign.

Para Elena y Clara.

Contents

List of Figures

List of Tables

Acknowledgments

Coauthor information: Chapters 1 and 2 of this dissertation are coauthored with Rodrigo Carril at Universitat Pompeau Fabra, and Michael S. Walker at the Cost Assessment and Program Evaluation Unit of the U.S. Department of Defense. Chapter 3 is coauthored with Mushfiq Mobarak at Yale University. I am extremely grateful for the opportunity to work with them, for all the hard work they have put into our projects, and for all they have taught me along the way.

Acknowledgements: I am extremely grateful to my advisor and mentor, Steve Tadelis, for his enormous support, guidance, and dedication throughout these years at Berkeley. He has played a crucial role in motivating me to tackle ambitious and economically meaningful research questions, and was always available to provide support with a kind word and wise advice above and beyond academic life. I am also profoundly in debt to Kei Kawai, who patiently encouraged me to think carefully about the theoretical and econometric models. I want to thank Ben Handel and Reed Walker for their fantastic advice and support during my time at Berkeley. Their guidance has been fundamental for my training as an applied economist. I also want to thank the generous support and encouragement I received from Ernesto Dal Bo, Jonathan Kolstad, Ricardo Perez-Truglia, Demian Pouzo, Chris Walters, Guo Xu.

This dissertation would not have been possible without financial support from the Haas School of Business. The project discussed in the third chapter was only possible due to Sernapesca's cooperation and JPAL-LAC's research and data collection support. It was funded by the Macmillan Center at Yale, the Center for Business and Environment at Yale, and Sernapesca. The first two chapters received funding from the Institute for Business Innovation at Haas.

The last six years in Berkeley have been wonderful thanks to many friends and colleagues who helped make it my home. There is a special place in my heart for Dario Tortarolo, Pablo Muñoz, Mathieu Pedemonte, Roman David Zarate, Chris Campos, and Ezequiel Garcia-Lemberman. Their friendship and support during these years was invaluable. Our discussions about research and life have made me grown as an economist and as a human being. I also want to thank enormously Yasir Khan, Gauri Subramani, Thiago Scot, Cristobal Otero, Damian Vergara, Julian Langer, and Cristian Ugarte. They made this journey very enjoyable. I also keep great memories of my friends and teammates of the different soccer teams I played; *Garra Latina*, *Smoke Sellers*, *Invisible Hand*, and *Goal Diggers*. Finally, but not least, life at Berkeley was great because of Malena Arcidiacono, Isabel Hincapie, Josefina Edwards, Matias Alcalde, Isidora Aldunate, and Monica Zegers.

I am especially grateful to my family, who has always been there for me. My parents, Nacha and Andres, have supported me in every decision that I have made. Their effort and sacrifices made this step possible. I also want to thank my parents-in-law, Elena and Jaime. They have adopted me as one of them, and I will always be grateful for their love and support in all these years. Most important of all is my wife, Elena, who has been with me every step of the way. I could not have made this without her love, patience, and support. She is the pillar of my life. Finally, Clara reminded me, with her smile and joy, that there are more important things in life than work and research.

Chapter 1

Competition under Incomplete Contracts and the Design of Procurement Policies I: Effects of Publicity

1.1 Introduction

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, expanding competition for contracts that involve customized obligations and deliverables could allow under-qualified contractors to win, leading to deficient execution ex-post. Therefore, the assessment of competitive mechanisms in procurement should account both for potential benefits due to price reductions, as well as for potential adverse effects due to poor execution.

An empirical investigation of this trade-off is complicated, in part due to the need for comprehensive data on contract execution and a compelling research design. In this paper, we aim to make progress on both of these fronts to study the equilibrium effects of enhancing competition for procurement contracts in acquisition price and execution quality. We focus on U.S. Department of Defense (DOD) procurement, a setting of independent interest given that it awards $500 billion in procurement contracts per year, representing a sizable fraction of the U.S. economy. Moreover, this setting provides us with policy variation in the degree of contract competition, as well as with detailed administrative data throughout the life-cycle of each DOD contract, from design through to 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.[1] We exploit the discontinuous nature of these requirements to estimate the effect of enhanced contract publicity on four sets of outcomes: (i) the level of competition for the award, (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.

We start by analyzing the price effects of contract publicity. We do this by investigating the observed contract price densities of publicized and non-publicized contracts.[2] We then estimate the effects of publicizing contract opportunities on three sets of non-price outcomes: the level of competition, the characteristics of the selected vendors, and post-award performance. We do this by implementing a Regression Discontinuity Design (RDD) for contracts with an expected award amount close to \$25,000. The discontinuity in publicity requirements at this threshold generates a convenient natural experiment for studying the impact of the policy on these non-price outcomes.

We find that contract awards advertised through the government platform see an increase in the number of bids of roughly 60%, 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% less likely to be small businesses, and are located 60% farther from the buying agency. Furthermore, we find that increased competition leads to contract price reductions: publicized contracts are, on average, awarded at 6% lower prices. However, advertised contracts result in worse ex-post performance: the probability of experiencing cost-overruns and delays in the implementation stage increase by 7% and 8%, respectively. The latter results are driven by service contracts—as opposed to goods purchases—and by contracts that we ex-ante characterize as more complex. These results are robust to different estimation approaches and there is little evidence of buyers bunching at these advertising thresholds that could lead to bias in the RDD results.

Taken together, our 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. We benchmark price reductions ex-ante with increases in cost-overruns ex-post, and we find that intensifying competition only reduces overall contract costs for "simple" product categories, while increasing overall costs for relatively complex contracts. Overall, we find that suppliers' identity matters for explaining the variation in contract outcomes. Promoting competition hinders the buyer's ability to restrict participation to

[1] From a policy perspective, the volume of contracts impacted by this regulation make its implications economically meaningful. In 2018 alone, the DOD publicized contract solicitations valued at \$ 5.56 billion dollars via the online platform `FedBizzOpps.gov`.

[2] Our method is robust to the existence of strategic bunching below the threshold, aimed at avoiding publicizing certain contracts. In fact, we separately quantify the extent of strategic bunching and the price effects of publicizing contracts.

qualified vendors, while attracting the participation of vendors who tend to perform poorly ex-post.[3]

This paper contributes contributes to the literature examining transactions under incomplete contracting (Williamson, 1976; Goldberg, 1977; Hart and Moore, 1988).[4] This literature is mostly theoretical, with only a small number of empirical papers studying the interplay between competitive mechanisms and contract outcomes (Spulber, 1990; Bajari, McMillan, and Tadelis, 2009; Decarolis, 2014). Our paper departs from existing work along two relevant margins: first, existing papers focus on different award mechanisms (e.g., auctions versus negotiations), while our empirical framework keeps the award mechanism fixed but exploits variation in the set of relevant potential sellers; second, unlike existing literature that concentrates on studies of specific product categories, our sample includes a wide range of product categories purchased by the DOD. This allows us to provide a comprehensive picture of the implications of promoting competition, arriving at different conclusions for different types of purchases.

We evaluate the effects of publicity on award prices, using an empirical framework that builds upon existing papers that use density analysis to achieve nonparametric identification of behavioral parameters (Saez, 2010; Kleven and Waseem, 2013; Kleven, 2016; Chernozhukov, Fernandez-Val, and Melly, 2013). Our methods contribute to the existing papers by proposing a novel empirical framework that separately identifies the existence (and extent) of manipulation in the assignment variable—i.e. bunching—from treatment effects.[5]

The rest of the paper is organized as follows. Section 1.2 provides background on the U.S. DOD procurement system and the data we use for our analysis. In Section 1.3, we provide evidence on the effects of contract publicity on a range of relevant outcomes. Section 1.4 concludes.

1.2 Setting and Data

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 more than 3,000 different contracting offices that are part of an executive or independent agency.[6] The workforce in charge of public contracting is made up of over 35,000 contracting

[3]The main alternative hypothesis is that vendors behave differently ex-post depending on the level of competition ex-ante. We explicitly test and reject this *moral hazard* explanation in favor of our *adverse selection* hypothesis.

[4]More recent developments include Laffont and Tirole (1990); Tirole (1999); Chakraborty, Khalil, and Lawarree (2020)

[5]Our methods extend recent developments by (Jales and Yu, 2017) emphasizing that density distributions around policy thresholds can be used to identify the effects of policy.

[6]Executive agencies are headed by a Cabinet secretary, like the Department of Defense, the Department of State, or the Department of Health and Human Services. Independent agencies, which are not part of the

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

FAR Part 5 requires publicizing contract opportunities in order to "increase competition", "broaden industry participation", and "assist small businesses (and other minority businesses) in obtaining contracts". Since October 1, 2001, contract actions that exceed \$25,000 must be publicized in an online government-wide platform which we will refer to as FedBizOpps (or FBO).[7] This implies uploading a request for quotes with a full description of the good or service being requested, and the instructions to submit the bids. We will refer to this synopsis document as a contract *solicitation*. Most of the contracts in this dollar range are awarded to the lowest price quote that is technically acceptable according to the specifications.

Officers with contracts that are not expected to exceed this threshold are not required to publicize in FedBizOpps; however, they are still free to use it if they want to increase contract visibility.[8] The regulation 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 publicized contracts around the threshold, we anticipate that compliance may be far from perfect, given the voluntary nature of the rule below this value and the availability of exceptions above. Appendix A.3 describes the details of the policy and the website.

Data

We use two complementary sources of data. The first consists of the historical files from FedBizOpps, which provides detailed information on pre-award notices (i.e. solicitations) posted on the platform. 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.

We merge awards from FPDS-NG to notices on FedBizOpps using the solicitation number. Note, however, that while FPDS-NG contains the universe of federal awards, FedBi-

Cabinet, include the Central Intelligence Agency, the Environmental Protection Agency, and the Federal Trade Commission.

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

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

zOpps only has the notices posted on the website. From this matching process, we construct a dummy variable that is equal to 1 if we are able to merge a contract with any pre-award notice on FedBizOpps, in which case we say the contract was *publicized*. Appendix Figure A.1.1 describes the typical timeline of events surrounding the life-cycle of a contract, and the appropriate data source that records that information.

In addition, we observe detailed information for each contract award, including the dollar value of the funds obligated, a four-digit code describing the product or service, codes for the agency, sub-agency, and contracting office making the purchase, the identity of the private vendor, the type of contract pricing, the extent of competition in the award, characteristics of the solicitation procedure, the number of offers received, and the applicability of a variety of laws and statutes. For additional details on the construction of the dataset, see Appendix A.3.

The analysis sample consists of all competitively awarded and definitive contracts[9] with award values between $ 10,000 and $ 40,000, awarded in fiscal years 2015 through 2019 by the Department of Defense (DOD), for products and services other than Research and Development (R&D).[10] Table 1.1 presents summary statistics of the sample. In total, there are roughly 86,000 contracts awarded by 597 contracting offices to almost 30,000 firms. Contract durations are expected to be 54 days on average and 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.[11] The Department of the Navy and the Army each account for more than 40% of the contracts, with the rest being mostly awarded by the Air Force. Winning vendors are often geographically close to the contracting offices, with both located in the same state in 2 out of every 3 contracts. Finally, 75% of suppliers are characterized as small businesses.

We also observe rich information about the type of good and service that is contracted upon. Each award is classified into one of 1,479 possible standardized 4-digit alphanumeric codes. These can be aggregated into 101 broader 2-digit product categories, 77 goods and 24 services. Table 1.2 shows the top 10 most common 2-digit good and service categories. The most common product categories are ADP Equipment Software, Medical Equipment and Supplies and Maintenance and Repair Equipment.

[9]Federal contracts can be broadly categorized into two types: definitive contracts (DCs) and indefinite delivery vehicles (IDVs). DCs are stand-alone, one-time agreements with a single vendor for the purchase of goods or services under specified terms and conditions. See Carril (2019) 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.

[10]The Department of Defense represents 58% of overall federal spending and more than 60% in the restricted sample. We exclude R&D awards because are subject to a unique set of acquisition rules, see FAR Part 35.

[11]More than half of the awards are set aside for a particular type of firm (typically, small business). Set-asides are a major factor of acquisition strategy in the DOD; 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.

1.3 The Effect of Competition on Contract Outcomes

We study the effects of publicly soliciting procurement awards on contract outcomes. As described in Section 1.2, federal regulation introduces a publicity requirement at \$25,000. In this section, we propose an empirical framework to exploit this discontinuity using different approaches. Section 1.3 describes the key identification challenges in our context. In Section 1.3, we propose an empirical method to estimate the effects of publicly soliciting contracts on award prices, based on the observed distribution of awards. In Section 1.3, we study the impact of publicity on other contract outcomes using a Regression Discontinuity Design (RDD). These results will serve as the basis for the development of our model in Section 2.2. For ease of exposition, in the main text we present a simplified version of our framework. Appendix A.4 presents all formal results and a more complete description, including all implementation details.

Preliminaries

For each contract in our data, we observe *realized* award prices (p), along with agencies' decisions to publicly solicit a contract in *FedBizzOpps.gov* prior to its award (decision that we denote as $D \in \{0, 1\}$). Since soliciting decisions are not made at random, we leverage the variation introduced by the regulation, which discontinuously affects the likelihood of public solicitation at an arbitrary threshold ($\bar{p} = 25,000$), in terms of the contract's *expected* award price ($\tilde{p}$).

We face two fundamental empirical challenges. First, we do not observe ex-ante estimated prices $\tilde{p}$, but only ex-post realized prices p. Second, since contracting officers know the policy threshold, this may generate incentives to modify the purchase in a way that makes the ex-ante estimate fall below $\bar{p}$. This behavior would result in bunching on the distribution of ex-ante prices, generating an excess amount of contracts estimated to be at or slightly below $\tilde{p} = \bar{p}$.

We can see these aspects more explicitly by writing down the observed award price in terms of potential outcomes. Let $p^B(\tilde{p})$ be the price observed when the buyer chooses to bunch, conditional on an ex-ante estimate of $\tilde{p}$. Similarly, denote as $p^D(\tilde{p})$ the *potential* price that we would have observed for a contract, conditional on an ex-ante estimate of $\tilde{p}$ and a publicity decision of $D \in \{0, 1\}$. We can write the ex-post realized award prices as:

$$p = p^0(\tilde{p}) + D \cdot \left[p^1(\tilde{p}) - p^0(\tilde{p})\right] + B \cdot (1 - D) \cdot \left[p^B(\tilde{p}) - p^0(\tilde{p})\right] \tag{1.1}$$

In what follows, we propose a method to nonparametrically recover information about the distribution of $\tilde{p}$, as well as the distribution of the price effects of publicity $[p^1(\tilde{p}) - p^0(\tilde{p})]$, from the distribution of observed awards p and publicizing decisions D. Intuitively, the method hinges on the comparison between the observed empirical distributions of award prices and estimated counterfactual distributions that are stripped of the confounding influence of bunching and competitive price effects. We proceed in three steps. First, we lay out the assumptions we rely on. Second, conditional on these assumptions, we characterize the distribution of densities around the regulatory threshold. Finally, we propose an estimation strategy that recovers the counterfactual objects of interest.

Density Analysis: Recovering Expected Awards and Estimating Price Effects

Baseline Assumptions

Consider a series of observed contract awards $t \in \{1,..,T\}$. For convenience, we normalize award prices as log-deviations from the policy threshold. That is, $p_t^{norm} = \log(p_t) - \log(\bar{p}_t)$. Unless there is ambiguity, below we simply drop the *norm* superscript, so whenever we refer to an award price variable (either ex-ante or ex-post) these are measured in logs and centered around zero, which corresponds to the policy threshold. In our baseline approach, we rely on the following assumptions:

A1 Ex-ante estimated prices are independently drawn from a smooth distribution, $\tilde{p}_t \overset{\text{iid}}{\sim} F_{\tilde{p}}(\cdot)$ with density $f_{\tilde{p}}(\cdot)$.

A2 Ex-ante estimated prices are equal to the price that the buyer would obtain without publicizing and absent any bunching, i.e., $p_t^0(\tilde{p}_t) = \tilde{p}_t$.[12]

A3 Relative to soliciting privately, publicizing solicitations leads to a (log-)linear random price effect. In particular, $p_t^1(\tilde{p}_t) = p_t^0(\tilde{p}_t) - \gamma_t$, with $\gamma_t \sim F_\gamma(\cdot)$, and $\gamma_t \perp \tilde{p}_t$.

A4 The probability that a contract is publicly solicited is a continuous function of its ex-ante estimated price, except at the threshold, where this probability jumps discontinuously. That is, $\Pr(D_t = 1|\tilde{p}_t) \equiv \tilde{\pi}_D(\tilde{p}_t) = \tilde{\pi}_D^*(\tilde{p}_t) + \delta \cdot \mathbf{1}[\tilde{p}_t > 0]$, for a continuous function $\tilde{\pi}_D^*(\cdot)$ and $\delta > 0$.

A5 Bunched contracts come from some region immediately above the threshold. That is, there exist $p_H > \bar{p}$ such that bunching is never chosen for $\tilde{p} > p_H$.

The first assumption is standard, requiring the latent density of ex-ante prices to be smooth. Importantly, we impose no additional parametric restrictions on $f_{\tilde{p}}(\cdot)$. Our second assumption is substantially stronger, requiring that buyers perfectly anticipate the price that they would obtain from a non-publicized solicitation. We adopt this for simplicity in our baseline approach. In the general exposition of Appendix A.4 we allow non-publicized contract awards to be anticipated with error, and show in Appendices A.4 and A.4 how to account for this in our key estimates.

Assumption **A3** imposes a (log-) linear price effect, i.e., constant *proportional* expected price effect in dollar terms.[13] Importantly, this is true *in expectation*: we do allow different contracts with the exact same estimate to lead to potentially very different realized awards,

[12] In other words, $\tilde{p}_t$ represents a perfectly accurate estimate of the price of soliciting without online posting.
[13] In dollar terms, the assumption is that $\log p - \log \bar{p} = \log \tilde{p} - \log \bar{p} - \gamma$, so that $p = \tilde{p} \cdot \exp(-\gamma)$.

because γ is a random variable. Since, in practice, our method is implemented on a relatively narrow window of award values (roughly within \$15,000 of the threshold in the main specifications), we view this assumption as not overly restrictive.

The fourth assumption requires continuity of the (unobserved) probability of treatment given ex-ante price estimates, except at the threshold where the policy generates a discrete positive jump. Our later results in Section 1.3, where we find a strong first stage jump in an otherwise continuous $\Pr(D_t = 1|p_t)$ function, are consistent with this assumption. Finally, **A5** states that bunching responses are somewhat local to the threshold from above. This region may not be small, but it is important that there exists a sufficiently large value of awards such that we do not see any bunching responses above this level. The assumption that these types of responses are confined to a window around the threshold is standard in the bunching literature and realistic in our setting.

Characterizing Award Densities

First, consider a baseline scenario in which there's no possibility to bunch, and there are no price effects of publicity. That is, suppose that $\gamma_t = 0$ and that $B_t = 0$ for all contracts. Our assumptions imply that both p_t^0 and p_t^1 will be exactly equal to the price estimate $\tilde{p}_t$. This will lead to distributions of publicized and non-publicized awards, such as the ones depicted in Figure 1.1. On the one hand, the total observed density of contract prices equals the density of ex-ante prices $f(\tilde{p})$, which is smooth. On the other hand, the regulation introduces a jump in the fraction of publicized solicitations at the threshold. Because there are no price-effects of publicity, the jump in the density of advertised contracts is compensated one-to-one with a drop in the density of non-publicized contracts.

Now suppose that we allow buyers to bunch: buyers can modify the purchase characteristics to reduce the estimated price and circumvent the regulation. This will translate into a non-publicized award of $p_t^B = \bar{p}$. Figure 1.2 (a) and (b) respectively show the effects of bunching on the densities of non-publicized and publicized contracts. Bunching generates an excess mass of *non-publicized* awards right below the threshold, as shown in the left panel. We do not impose assumptions on what would have been the treatment status of these contracts had bunching been forbidden. Therefore, the excess mass below the threshold in the non-publicized density has to be equal to the sum of the missing masses to the right of the threshold in the non-publicized and publicized graphs. Finally, note that, because of **A5**, there is a point to the right of the threshold where bunching responses no longer occur, as seen by the fact that the solid and dashed lines converge in Figure 1.2.

We consider next the role of price effects on top of bunching responses. Figure 1.2 (c) and (d) respectively show these effects for the density of non-publicized and publicized awards. Because price effects only affect advertised contracts, the mass in panel (c) is unaffected by γ_t. On the contrary, the full distribution of publicized awards is "shifted" by price effects. We can break this down into two margins. First, the distribution is shifted to the left based on the size of $E[\gamma]$, since publicized awards are now, in expectation, equal to $\mathbb{E}[p_t^1|\tilde{p}_t] = \tilde{p}_t - E[\gamma_t]$.

Second, the sharp discontinuity is "smoothed-out", because of the variance in the individual price effects γ_t.

Our estimation method relies on "reverse-engineering" these effects to recover key latent magnitudes from observed awards. We can separately identify price effects from bunching responses, even though they both generate excess mass below the threshold in the award density. Intuitively, the key insight of our method is that the excess mass generated by bunching is manifested only in the distribution of non-publicized contracts, whereas price effects influence only the distribution of publicized contracts. Our identification argument, therefore, critically relies on our ability to observe separately the densities of "treated" ($D_t = 1$) and "control" units ($D_t = 0$).

Estimation Method and Main Results

Based on these ideas, and building upon bunching methods originally developed to measure behavioral responses to taxes (see Kleven, 2016), we explain our approach to nonparametrically recover price effects and ex-ante award prices from the observed distribution of contract awards. The discussion is brief and relies heavily on the intuition from the graphical analysis in Section 1.3. For most formal details, we refer to the discussion in Appendix A.4.

The estimation procedure consists of three steps. First, we estimate the average price effects of publicity, $E[\gamma_t]$. We then obtain estimates of the distributions of unobserved ex-ante prices for both publicized $\left(f_{\tilde{p}}(\tilde{p}_t|D_t = 1)\right)$ and non-publicized $\left(f_{\tilde{p}}(\tilde{p}_t|D_t = 0)\right)$ contracts. Finally, we impose an additional assumption that allows us to recover the full distribution of price effects $(F_\gamma(\gamma_t))$ nonparametrically.

Binning and Notation. We first discretize the range of normalized award values into a set of equally-sized and right-inclusive bins around the threshold $b \in \{-R, (-R+1), ..., -1, 0, 1, ..., (R-1), R\}$. Note that bin $b = 0$ includes awards right at, or slightly below, the policy threshold. Let $\{n_b^d\}_{b=-R}^R$ be the frequency distribution of observed contract prices conditional on treatment (publicity) status $D_t = d$, for $d \in \{0, 1\}$, so that n_b^d denotes the number of contracts with treatment status d and observed award value in bin b. Likewise, let $\{\tilde{n}_b^d\}_{b=-R}^R$ represent the (unobserved) frequency distribution of latent ex-ante prices.

Step 1: Average price effects. Our method builds upon the fact that, relative to ex-ante prices, the mean and the variance of linear price effects impact the distribution of publicized contracts in distinctive ways. Intuitively, if $E[\gamma_t] = \bar{\gamma} = 0$, then the difference between expected and observed prices would only stem from the stochastic nature of γ_t, and would manifest itself only close to the threshold.

In particular, a mean-zero γ_t implies that a publicized contract with ex-ante price right at the discontinuity has an equal probability of being observed above or below the threshold. Thus, relative to the counterfactual distribution, the excess of contract mass below the threshold —due to contracts that "fall" from above— equates to the missing mass of contracts missing above the threshold. Conversely, if the average price effect $E[\gamma_t] = \bar{\gamma}$ was greater (smaller) than zero, most contracts with an ex-ante price exactly at the discontinuity would be shifted to the left (right), generating an excess of mass below the threshold that would (not) exceed the missing mass above. This intuition is shown in Appendix Figure A.1.2. The key conclusion is that the integration constraint equating missing and excess mass is only met when there is a mean price effect of zero.

The estimation of $\hat{\gamma}$ builds on this logic. We iterate over possible guesses of $\widehat{\tilde{\gamma}}$, shifting the distribution of publicized contracts accordingly, until we find a value such that the integration constraint is satisfied. The excess and missing masses are measured relative to an estimated counterfactual distribution $\{\widehat{\tilde{n}_b}\}$, which is obtained using a standard bunching estimation procedure: i.e., we interpolate from outside an excluded window around the threshold, using a polynomial fit. Figure 1.3 shows the implementation of our method. In panel (a) we show the empirical distribution of non-publicized contracts $\{n_b^0\}$. In panel (b) we present the empirical distribution of publicized contracts $\{n_b^1\}$, as well as a *shifted* distribution of publicized contracts $\{n_b^{1,s}(\widehat{\tilde{\gamma}})\}$, for $\widehat{\tilde{\gamma}} = 0.0595$ (standard error of 0.0201, computed via bootstrap), which is our main estimate of price effects. Finally, Panel (c) shows the distribution $\{n_b^0 + n_b^{1,s}(\widehat{\tilde{\gamma}})\}$, as well as the polynomial fit of the bunching estimation, which corresponds to $\{\widehat{\tilde{n}_b}\}$. We therefore estimate that, on average, publicizing a contract leads to a reduction in award price of 0.06 log-points, equivalent to \$ 1,456 at the discontinuity.

Step 2: Distributions of ex-ante prices. Having obtained an estimate of mean price effects $\widehat{\tilde{\gamma}}$, as well as a counterfactual distribution of latent ex-ante awards $\{\widehat{\tilde{n}_b}\}$, we now recover separate counterfactual distributions of ex-ante prices for both publicized and non-publicized contracts. These distributions should be continuous everywhere, *except* at the threshold, since the policy generates a discontinuous jump in advertised contracts mirrored by a discontinuous dip in non-publicized contracts. We estimate the magnitude of this jump (which we denote Δ) using the same logic as before: first undo the jump by shifting (this time, vertically) the right part of the distributions of $\{n_b^0\}_{b=-R}^R$ and $\{n_b^{1,s}(\widehat{\tilde{\gamma}_b})\}_{b=-R}^R$, and then eliminate the effect from bunching responses and the variance in price effects by interpolating values within a window around the threshold. The procedure searches for the magnitude that maximizes the fit of the interpolations.[14] The logic is that, if the vertical shift is too low or too high, the polynomial interpolation will fit poorly just outside of the excluded window.

Figure 1.4 shows the results. Our procedure estimates the counterfactual distribution of both publicized and non-publicized contracts, stripped down from price effects and strategic bunching responses. From the distribution of non-publicized 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% of the value of the density at the threshold. This magnitude will be crucial to account for the effects of this manipulation on our RDD estimates in Section 1.3. However, we can already infer that, since the extent of bunching is arguably modest, its impact on our estimates will be limited as well.

Finally, in Panel (b), we recover the sharply discontinuous distribution of publicized awards that would be observed if $\gamma_t = 0$ for all t. Our ability to observe distributions with

[14] In particular, we minimize the average between the sum of squared residuals for the treated (publicized) and control (non-publicized) distributions.

and without price effects is the key to recover the full CDF of price effects, as described in our last step.

Step 3. Full distribution of price effects. With one additional assumption, we can nonparametrically recover the full CDF of price effects F_γ, given our estimates of latent ex-ante price distributions and our observation of realized price distributions. In Appendix A.4, we show that:

$$F_{\gamma'}(x) \approx \frac{n^{1,s}_{b_x}(\bar{\gamma}) - \tilde{n}^1_{b_x}}{\Delta}, \qquad \text{for } x \in b_x, b_x \leq 0 \tag{1.2}$$

where $F_{\gamma'}(x)$ is the CDF of $\gamma_t - \bar{\gamma}$, so that it is re-centered around zero.

This equation tells us that we can recover part of the CDF of price effects from the difference between the (appropriately shifted) observed distribution of prices for publicized contracts and the latent distribution of ex-ante prices for this same group, scaled by the size of the discontinuity at the threshold. Intuitively, at a given value x below the threshold, the observed excess mass generated by the variance in individual price effects, is proportional to the probability that a contract with ex-ante price above the threshold is "moved down" to x by the price effect. The presence of a discontinuity Δ is key for this identification argument, which along with variation in x allows us to estimate $F_\gamma(x)$ for $x \leq 0$. This gives us the lower half of the CDF. Imposing symmetry of $F_\gamma(x)$ around its mean, we can recover the full CDF F_γ.

A6 The CDF of price effects $F_\gamma(\cdot)$ is symmetric, i.e. $F_\gamma(-x) = -F_\gamma(x)$.

Figure 1.5 depicts our nonparametric estimate of F_γ, along with a local polynomial smoothing. By construction, $F_\gamma(x)$ is symmetric and centered around our estimate of mean price effects $\widehat{\bar{\gamma}} = 0.06$. We see that publicity reduces award prices for 83% of the contracts.

Regression Discontinuity Design: Estimating Effects on Non-price 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 level of competition, 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:

$$Y_t = \alpha + \beta \cdot D_t + g(\tilde{p}_i) + X_t'\delta + \epsilon_t \quad , \tag{1.3}$$

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

$$D_t = \lambda + \gamma \cdot \mathbf{1}[\tilde{p}_t > \bar{p}] + g(\tilde{p}_t) + X_t'\eta + \nu_t \quad , \tag{1.4}$$

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

$$Y_t = \mu + \phi \cdot \mathbf{1}[\tilde{p}_t > \bar{p}] + g(\tilde{p}_t) + X_t'\kappa + \xi_t \quad . \tag{1.5}$$

Again, our key challenge is that we observe p_t, but not $\tilde{p}_t$ directly, and that the mapping between these variables is affected by price effects and possible bunching responses. However, with the procedure discussed above we can recover significant information about the latent distribution of $\tilde{p}_t$, which allows us to successfully exploit the discontinuity for estimating causal effects.

Consider first a naive RDD, described by versions of Equation 1.3, Equation 1.4, and Equation 1.5, where we simply replace ex-ante prices $\tilde{p}_t$ by realized observed prices p_t. The estimates obtained from this naive RDD 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, to transparently show how these elements affect the estimation.

In Appendix Section A.4, 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 modelling assumptions, we can write the conditional expectation of contract outcomes given *observed* prices $E\left[Y_t|p_t\right]$ 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.[15]

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 (non-publicized 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 in terms of how selection

[15] We also show in Appendix Section A.4 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 $p^0(\tilde{p})$. That is, we show that it is possible to recover causal parameters even if we drop assumption **A2**.

can influence RDD estimates. In Appendix Section A.4, we explain in detail how to derive these bounds in our setting, and how to calculate them using our estimate of excess bunching obtained in our density analysis.

Effects on Non-Price Outcomes

Naive RDD Results: We start with the naive RDD results, and then sequentially apply corrections to account for the specific issues present in our setting. We estimate specifications Equation 1.3, Equation 1.4, and Equation 1.5, 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 also present results from the robust local polynomial approach proposed by Calonico, Cattaneo, and Titiunik (2014). We present these naive RDD results visually, by plotting binned scatters of Equation 1.4 and Equation 1.5. In the next section we explicitly assess how these baseline estimates change as we consider the impact of price effects and (or) bunching responses.[16]

The results for the first stage Equation 1.4 are presented graphically in Figure 1.6. We see that the use of FedBizOpps jumps sharply past the $25,000 threshold of award amounts. The share of contracts that are publicly solicited in the government platform increases from roughly 30% at or slightly below $25,000, to 50% right above this threshold.

The reduced form specifications (Equation 1.5) are estimated on three sets of outcomes: the intensity of competition, winning vendor characteristics (including its relationship with the awarding office), and post-award performance. Most of the existing literature has studied these variables independently.[17] By studying them jointly, we can generate a comprehensive understanding of the mechanisms and implications of policies oriented to enhance competition.

Figure 1.7 shows how posting solicitations on FedBizOpps impacts the number of offers that a contract receives around the threshold. Contracts right above $25,000 (which are more likely to be publicly solicited), receive roughly 0.4 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 p.p.

These results indicate that encouraging the public posting of solicitations leads to the stated goal of increasing competition by attracting additional bids. 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 1.8 shows that this is not the case. In Panel (a), we see that publicized contracts are awarded to vendors that are relatively larger,

[16] Figure A.1.4 presents RDD plots for baseline variables. We find that baseline contract design characteristics are balanced around the threshold, with the exception of goods vs. services. There are more services right above the threshold. The difference is noisy and against possible selection patterns. All of our baseline estimates are robust to the inclusion of a service dummy as control.

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

as measured by a reduction of the probability of awarding the contract to a small firm.[18] 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 located geographically more distant from the contracting office location. These results suggest that marginal entrants attracted by the 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 (e.g. Decarolis, 2014; Kang and Miller, 2017; Decarolis et al., 2020; Carril, 2019). Because the data contain the total sum of payments and completion date expected at the time of the award for each contract, we can construct measures of cost overruns and delays by comparing these expectations to the realized payments and duration. These measures have been used by recent studies as performance proxies.[19]

Figure 1.9 presents the results. We find that the share of contracts with overruns and the share of contracts with delays increase by 2 p.p. and 1.5 p.p., respectively. These differences are statistically and economically 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. Figure A.1.5 in the Appendix 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 expected schedule. These results align with the findings presented in Figure 1.9: publicized contracts experience more problems during the execution stage.

Adjusted RDD Results: In this section, we present a series of refinements to our naive RDD results. First, we explore robustness to 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 1.3, 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). The price effect correction is related to existing methods that account for measurement error in the RDD framework (Pei and Shen, 2017). The key advantage of our setting is that we express the conditional expectation of contract outcomes given *observed* prices $E\left[Y_t|p_t\right]$, as an explicit function of magnitudes that we estimated in the density analysis, plus the causal parameters that we seek to recover. On the other hand, we follow Gerard, Rokkanen, and Rothe (2020)'s approach to account for the potential effect of bunching responses. They show that given an estimate of the extent of bunching —which

[18] The Small Business Administration (SBA) 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.

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

we obtained in Section 1.3—, we can bound the estimated treatment effects under "worst" and "best" case scenarios of how selection influences the RDD estimates. The width of these bounds shrinks as the extent of bunching decreases, converging to the point estimates in a case with no manipulation.

Table 1.4 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 et al. (2014)'s local polynomial estimates with robust bias-corrected standard errors. Overall, non-linear estimates are similar in magnitude and significance to simple OLS estimates. The third column present estimates that account for price effects in the treatment group (i.e. publicized contracts), following the method explained in Appendix A.4. The correction for price effects is relatively modest, and in most cases tends to amplify the naive results. This is consistent with the fact that the price effects smooth-out the discontinuity for the treatment group. Thus, under naive estimation, some publicized contracts are observed below the threshold when their original (ex-ante) price was above it.

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 do 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 similar magnitude than the worst-case upward bias introduced by bunching responses.

Taken together, these results imply that the strong visual evidence presented in Figures 1.6 through 1.9 is robust to fully accounting for the potentially confounding influence of price effects and strategic bunching by the buyer.

The Role of Contract Complexity

Our analysis includes a wide variety of transactions, from standardized goods to customized services. A procurement contract aims to regulate the nature of the expected transaction. Nevertheless, specifying possible contingencies is easier if the purchase involves a commodity-type product rather than an *ad-hoc* service. Thus, the more difficult it is to specify the need in a contract, the more variable will be the post-award performance. This explains why some product categories rarely experience performance issues ex-post, while others experience implementation issues in most of the contracts. Similarly, the effects of expanding competition on award prices are also likely to vary depending on the good or service's underlying complexity. For example, if bidders of relatively complex products are more heterogeneous in production costs, additional offers would lower contract prices more than when contractors are homogeneous. Thus, there are reasons to believe that the degree of contract complexity affects both prices ex-ante and prices-post.

To assess these mechanisms more directly, we leverage the rich heterogeneity of our data, which features 1,918 distinct product categories, and we proxy the degree of contract complexity based on the baseline level of post-award performance, which we define as the average cost-overrun experienced by all contracts below $20,000 for the same product category.[20] Table 1.5 describes the top and bottom product categories based on average delays and cost-overruns. Contracts for services experience substantially more issues in the implementation stage than goods. Using the measure of mean cost-overruns, we divide the contracts categories in our sample into quartiles of complexity, and re-estimate both price effects and RDDs on performance, separately for each of the four groups. We also consider the more simple heterogeneity of effects between goods and services.

Table 1.3 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 1.10 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% and of services by 7.8%. This effect corresponds to 4% for the least complex quartile, versus 9.6% for the top quartile of complexity.

The results are qualitatively similar for the impact of publicity on post-award performance. Figure 1.11 shows that the increase in overruns and delays that we reported in Figure 1.9 is driven by goods and services in the *top* quartile of complexity. We are unable to reject the null for the lower three quartiles. Overall, we see the effects in overruns outweighs the price reductions ex-ante for complex contracts. However, when the unit is "simple" the price reductions, although modest, exceed increases in cost-overruns ex-post.[21]

Evidence of Adverse Selection

Our results show that increasing the pool of bidders through publicity generates changes to contract prices and the subsequent contract execution. Overall, there are two classes of explanations through which we can rationalize the connection between publicity and contract outcomes.[22] The first explanation is that contractors modify their behavior depending on

[20] There are multiple ways of characterizing product complexity. We tried different approaches, e.g., using performance's standard deviation, indexing multiple variables, counting the number of words in the solicitation's description, etc. These classifications lead to roughly the same rank of products categories, and thus varying the definition does not threaten the general results. We use the mean of cost-overruns because it is transparent and easy to interpret. We get around the issue of classifying based on an outcome by focusing on contracts below $20,000.

[21] Figure A.1.13 in the Appendix shows RD plots for cost-overruns separating for 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.

[22] Publicizing contracts in FBO.gov impacts neither the contract's design nor the mechanism of selection. Appearing in FBO.gov solely affects the diffusion of information.

the publicity status of the contract (i.e. *moral hazard*).[23] The second explanation is that publicity allows the participation of suppliers that are "different," an that their performance capacity is unrelated to the identity of the buyer or contract's advertising (i.e. *adverse selection*).

We combine features of our setting with empirical methods used to detect asymmetric information (e.g., Chiappori and Salanie, 2002), in order to elucidate these mechanisms. In particular, we leverage that buyers often require the same product categories repeatedly over time; indeed, we observe multiple contracts for the same buyer-product combination, with variation in the size of the award and other dimensions. Moreover, from the supplier's 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 contractor's "types," relative to variation "within" contractor. To do this, we re-estimate the RDD analysis on post-award performance including contractor fixed-effects. The fixed effects demean contractors' performance, allowing us to test how much variation remains. In Appendix Table A.2.1, we show that the changes in performance disappear once we include contractor fixed effects. This implies that most of the variation on contract performance introduced by publicity is explained by variation across contractors, as opposed to "within" contractors.

Classifying Contractors' "Types"

Publicity increases the set of potential bidders, but seems to attract contractors who have limited capacity to execute contracts (as evidenced by there lower performance ex-post). We can extend our analysis by classifying contractors depending on their usual source of information about contract solicitations. Based on the patterns of contractor's participation, we identify two separate groups of firms: contractors who win awards without relying on publicity —which we refer to as *locals*—, and contractors that *only* win when contract solicitations are publicized —which we label *non-locals*. The logic is that, if a contractor wins without publicity, this indicates that the buyer informed her directly (e.g. through email or a phone call). The existence of direct communication reveals a buyer's preference for these contractors. Conversely, if a contractor requires a FedBizzOpps announcement to participate (and win), this suggests that there is no specific preference from that buyer for that contractor. This distinction came up frequently in conversations with procurement officers from several organizations.

To classify contractors empirically, we restrict the analysis to buyer-product combinations that are observed at least 4 times between 2013 and 2019, and which had at least one —but not all— contracts publicized.[24] Table 2.1 compares buyer-contractor distance and

[23]This could be rationalized if suppliers behaved differently depending on the buyer. For example, if a vendor receives contract information directly from the buyer, she could decide to absorb potential overruns to make sure she gets direct information again in the future.

[24]We noted that the Federal Procurement Data System (FPDS) sometimes missclassifies local buyers, assigning the same code to different branches that depend on a single (higher-level) office. This contrasts with the

performance for contracts performed by local and non-locals. The third column shows the mean difference of performance between these two groups. As a reference, if the information source is irrelevant, locals and non-locals would have similar outcomes. However, we observe that contracts executed by *non-local* contractors experience 16 percentage points (200%) more cost-overruns and 14.5 percentage points (110%) more delays than locals. This distinction will feature prominently in the empirical model of Section 2.2.

Discussion

This section provided evidence that 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. Leveraging detailed information about contracts' implementation, we also found that publicized contracts result in more cost-overruns and delays. Taken together, our results show that promoting contract competition involves a trade-off: it reduces contract award prices at the risk of awarding under-qualified contractors that are unable to execute contracts at the desired quality. Thus, the desirability of competition depends on which of these effects dominate. Furthermore, we find that this trade-off is heterogeneous, with both price effects and performance effects depending on the degree of contract complexity.

While this policy analysis is informative on the effects of promoting competition on contract outcomes, several questions about the underlying market structure that shape bidders' adverse selection remain unanswered. In particular, our reduced-form analysis does not allow us to evaluate equilibrium conditions under alternative policy designs that account for contract complexity and buyers' preferences. To make progress on these fronts, we now present an estimate an equilibrium model of competition promotion, firms' participation, and bidding decisions.

1.4 Conclusion

This paper studies the relationship between competition and procurement contract outcomes. Even though procurement contracts represent a key component of the economy, there is minimal evidence of the implications of policies oriented to expand competition, considering not only the award price but also the quality of the contract execution. We provide extensive evidence of the effects of enhancing competition through publicity, using the U.S. Department of Defense contracting market as a setting.

Our identification strategy leverages a regulation that generates quasi-experimental variation in the extent to which contract opportunities are broadly advertised to potential suppli-

nature of most procurement officers' job, who typically contract within a particular area, leveraging their local market knowledge. We address this misclassification by defining a buyer based on the office code *and* the Metropolitan Statistical Area (MSA) of the purchase. As before, the definition of a product category is given by the 4-digit PSC code.

ers. We find that contract publicity increases contract competition. The added competitive pressure results in lower acquisition prices; however, broader dissemination leads to a different pool of vendors, who perform worse ex-post. We further explore the implications of the key trade-off between price reductions ex-ante and worsen contract execution ex-post. Our analysis shows that the degree of contract complexity determines the scope of this trade-off. Promoting competition reduces contract costs only for simple transactions, as relatively complex ones are exposed to cost overruns and delays in the execution stage. This evidence highlights the role of contract incompleteness in the procurement setting. Moreover, the fact that contracts are incomplete introduces potential adverse effects to policies promoting competition. In chapter 2 we develop and estimate and model to study the primitive determinants of the estimated effects and how to design optimal competition policies accounting for incomplete contracts.

Figures

Figure 1.1: Award distributions in the absence of bunching and price effects

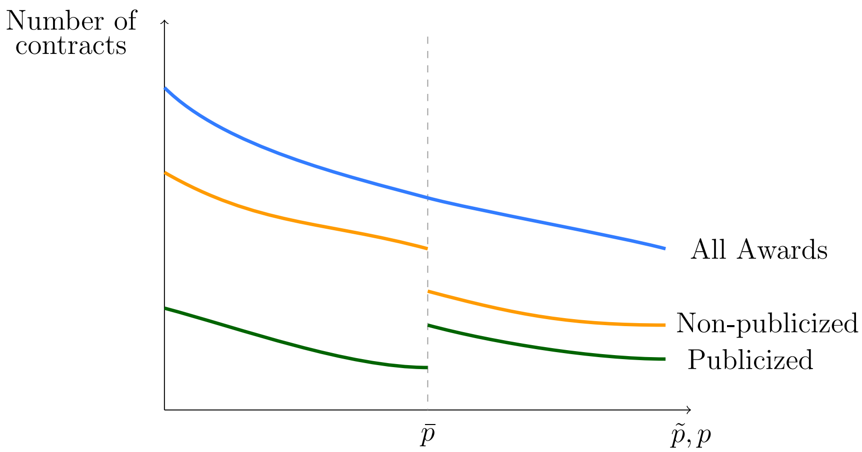

Notes: This figure shows an example of distributions of non-publicized and publicized awards, as well as the distribution of all awards (i.e. the sum of the two). The distribution of all awards is continuous by Assumption **A1**. The distributions of publicized and non-publicized contracts are discontinuous at the policy threshold $\bar{p}$. The jump in the number of publicized contracts is equal to the drop in the number of non-publicized ones.

Figure 1.2: Impact of Bunching and Price Effects on Award Distributions

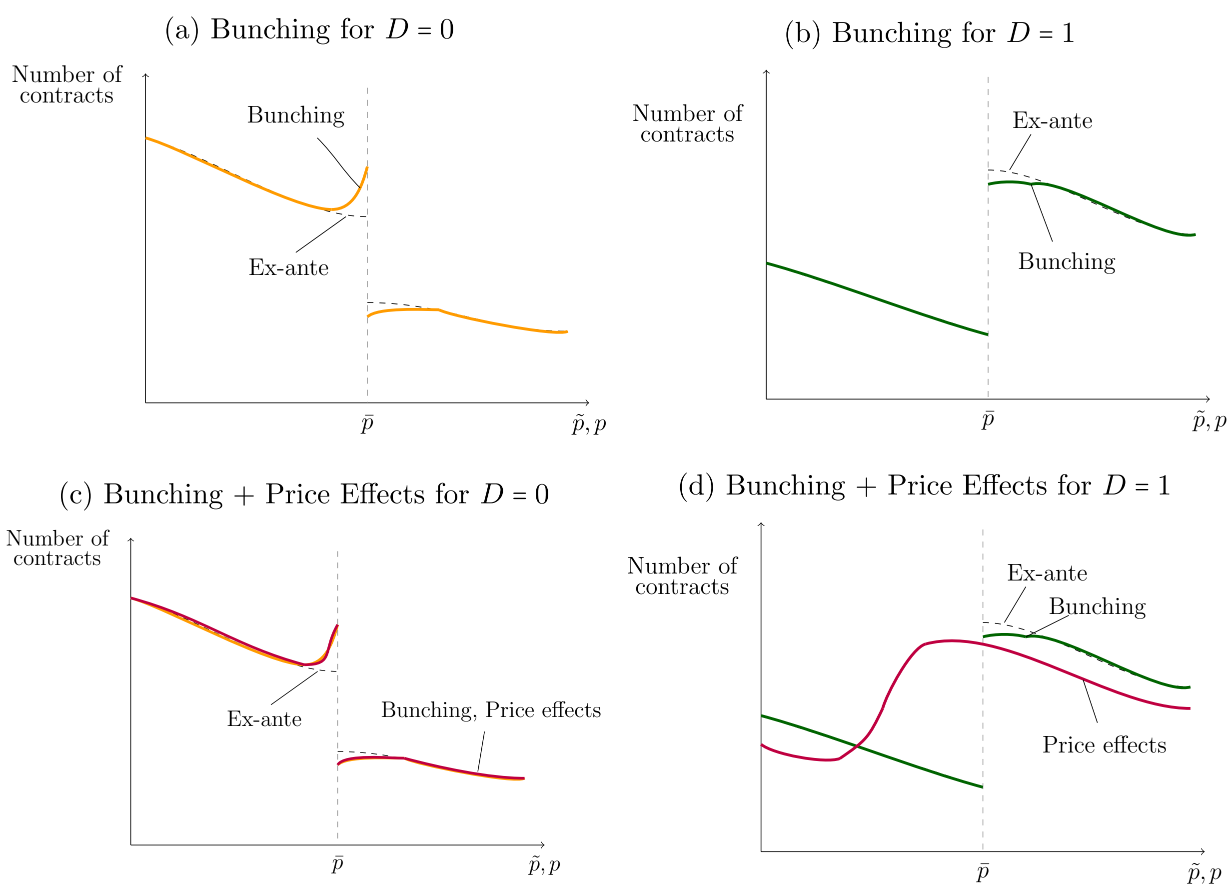

Notes: This figure shows conceptually how the distributions of non-publicized and publicized awards are impacted by the existence of both strategic bunching responses and price effects due to increased competition. Panels (a) and (b) show, respectively for non-publicized and publicized contracts, the distributions of ex-ante award prices ($\tilde{p}$, in dashed black lines), as well as realized award prices (p, in solid orange and green lines) when we allow for strategic bunching responses. Panels (c) and (d) plot the additional effect of having price effects associated with publicity (in solid red lines).

Figure 1.3: Estimating mean price effects and ex-ante prices

(a) Non-publicized contracts ($D = 0$)

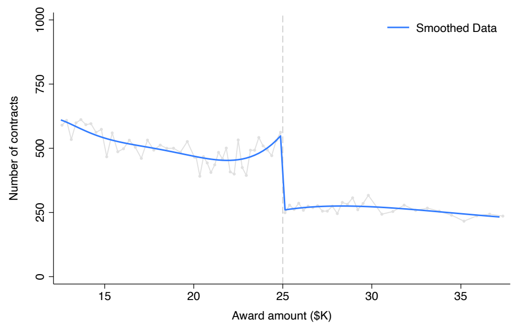

(b) Publicized contracts ($D = 1$)

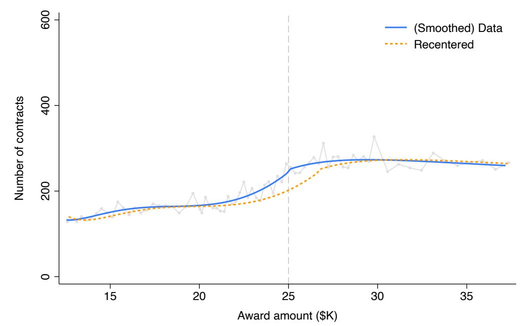

(c) All contracts

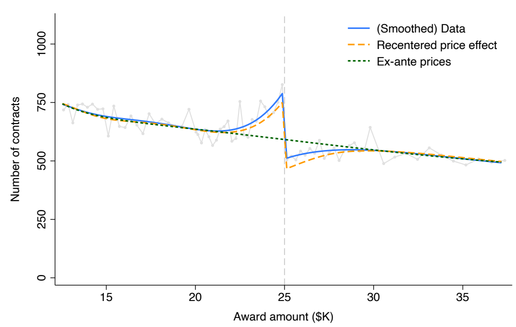

Notes: This figure shows the empirical distribution of the number of contracts at different price bins. Panel (a) shows the distribution of non-publicized contracts ($D = 0$). Panel (b) shows the distribution of publicized contracts ($D = 1$). Panel (c) displays the overall distribution, i.e., the sum of publicized and non-publicized contracts at every price. The blue line corresponds to a polynomial fit of degree five. The orange dashed lines in panels (b) and (c) represent the distribution of contract prices after re-centering publicized contracts by their price effect. The green dashed line in panel (c) represents the corresponding overall interpolation in the absence of price effects and bunching.

Figure 1.4: Estimating ex-ante prices

(a) Non-publicized contracts ($D = 0$) (b) Publicized contracts ($D = 1$)

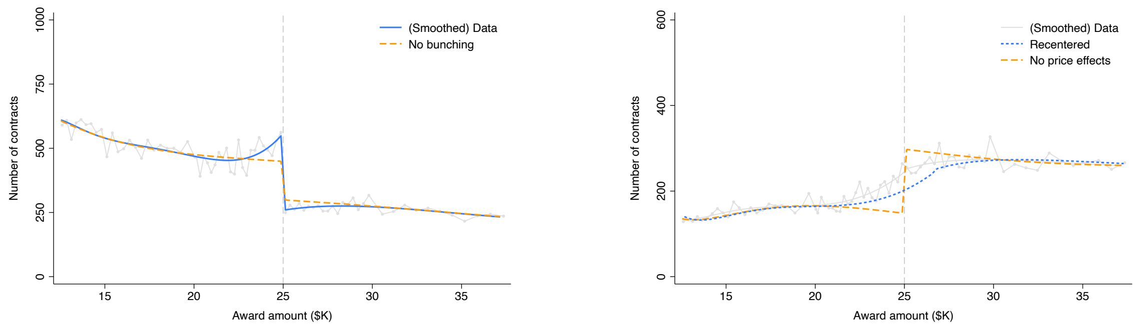

Notes: This figure shows the empirical distribution of the number of contracts at different price bins. Panel (a) shows the distribution of non-publicized contracts ($D = 0$). Panel (b) displays the distribution of publicized contracts ($D = 1$). The blue line corresponds to a polynomial fit of degree five. The orange dashed lines in panels (b) and (c) represent the counterfactual distributions in the absence of price effects and bunching. The counterfactual distributions stem from the proposed framework. In panel (a), The comparison between the solid blue and the dashed orange lines provide a visual interpretation of the mass of bunched contracts. The comparison between the dashed blue and the dashed orange lines in panel (b) inform visually about the distribution of price effects.

Figure 1.5: Distribution of Price Effect

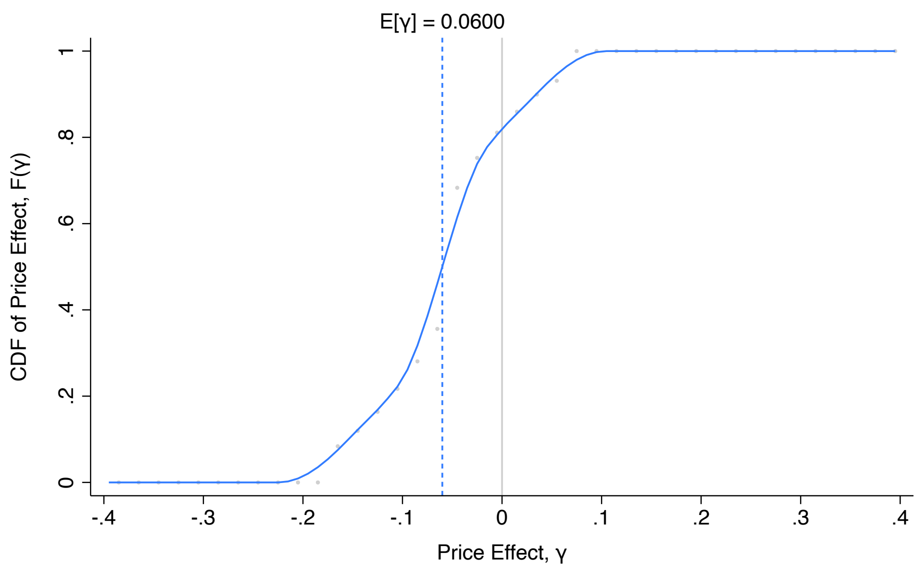

Notes: This figure presents the estimated CDF of the price effect parameters γ. The gray dots show actual point estimates given a discretization of the support of γ. The blue line corresponds to a kernel fit. This estimation procedure builds upon comparing the empirical densities to a counterfactual distribution of publicized contracts assuming no price effect. The counterfactual distribution is generated from the interpolation of a polynomial of degree 5. The dashed vertical line corresponds to the estimated mean effect.

Figure 1.6: Publicizing requirement and use of FedBizOpps

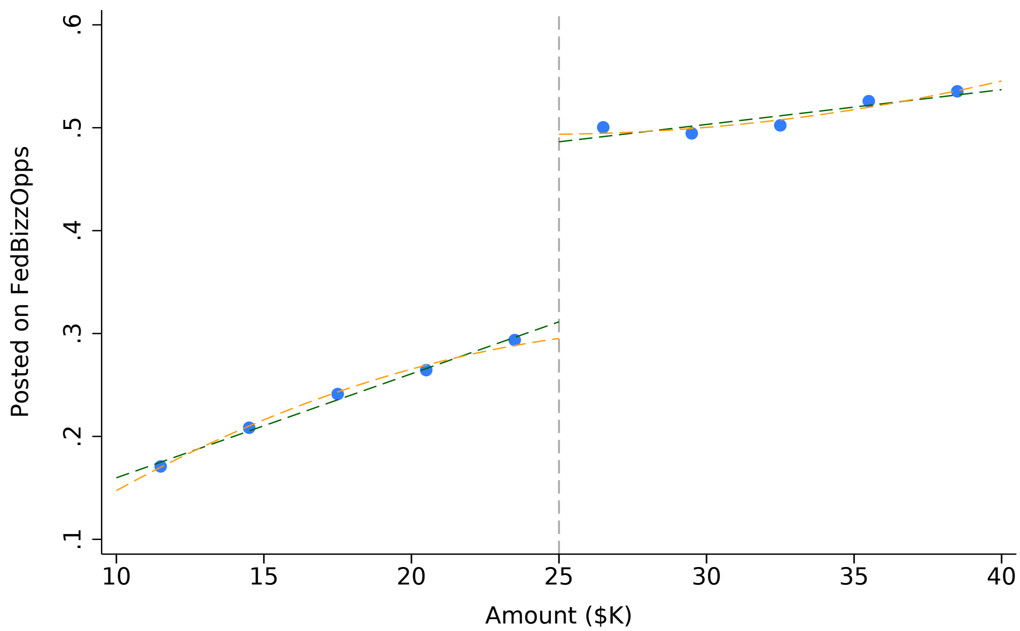

Notes: This figure shows the fraction of contracts posted on FedBizOpps by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data sources are FBO.gov and the Federal Procurement Data System-Next Generation. The sample consists of competitive, non-R&D, definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure 1.7: Publicity and intensity of competition

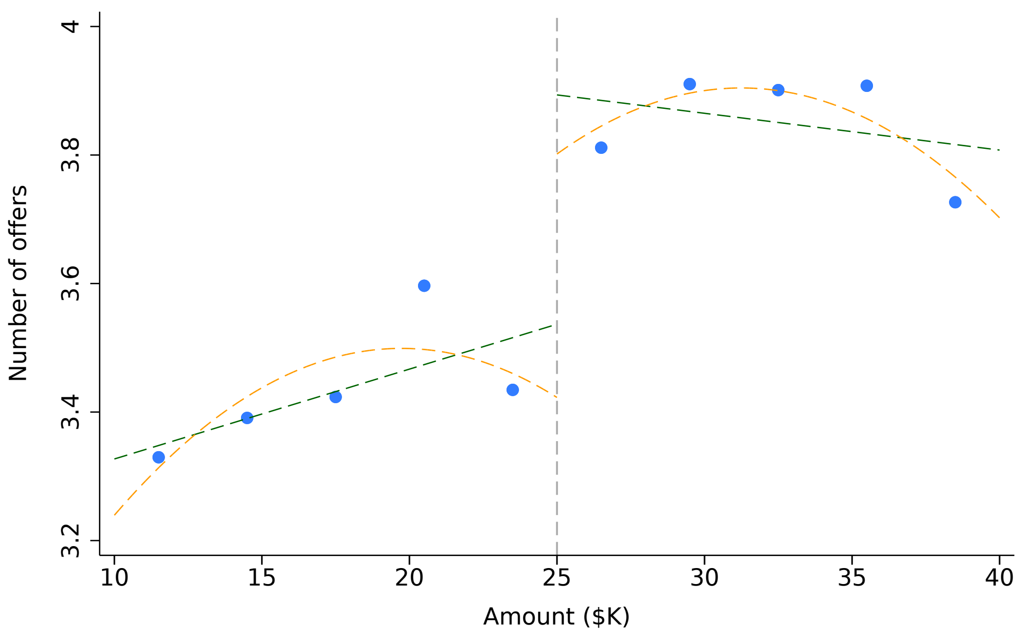

Notes: This figure presents four binned scatter plots, which depict the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data comes from FBO.gov and the Federal Procurement Data System-Next Generation. The sample consists of competitive, non-R&D, definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure 1.8: Publicity and the characteristics of the winning firm

(a) Contractor is a small business

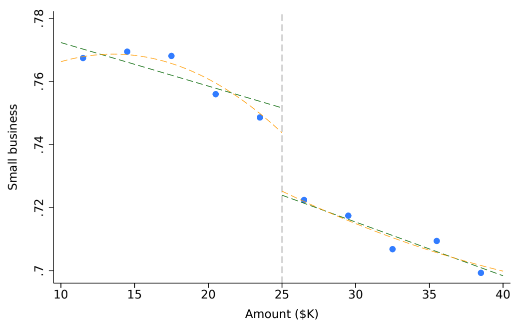

(b) Foreign firm

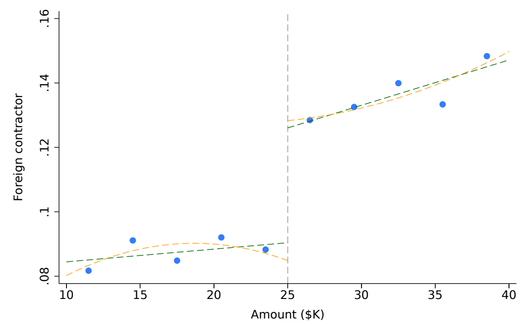

(c) Distance to the office

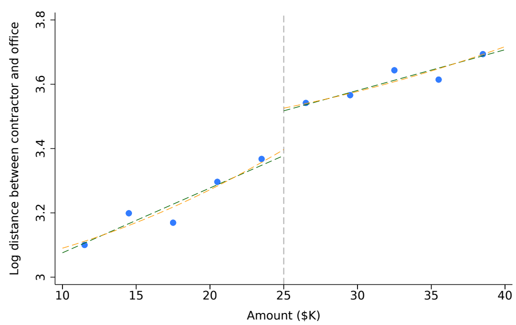

Notes: This figure presents four binned scatter plots, which depict an average outcome by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The outcome in each Panel is as follows: (a) indicator equal to one if the awarded contractor is a small business (based on SBA); (b) an indicator equal to one if the contract is awarded to a foreign vendor; (c) the natural logarithm of the distance (in miles) from the contracting office's location and the vendor location. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D, definitive, and competitively awarded contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure 1.9: Publicity and post-award contract performance

(a) Share of contracts experiencing cost-overruns (b) Share of delayed contracts

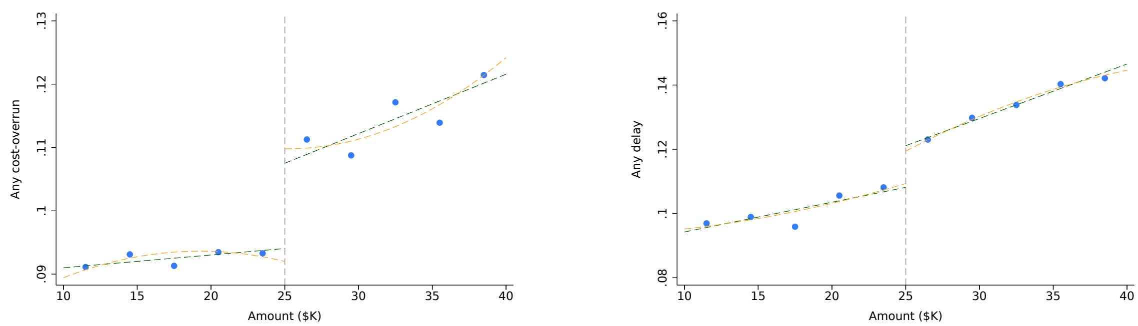

Notes: This figure presents four binned scatter plots, which depict an average outcome by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The outcome in each Panel is as follows: (a) the share of contracts that the actual obligated contract dollars exceed expected total obligations at the time of the award (i.e., cost-overruns); (b) the share of contracts whose actual days of contract duration exceed the expected days of duration at the time of the award (i.e., delays). The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D, definitive, and competitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure 1.10: CDF Price Effect by Sub-group

(a) Good vs. Service (b) Quartile of Complexity

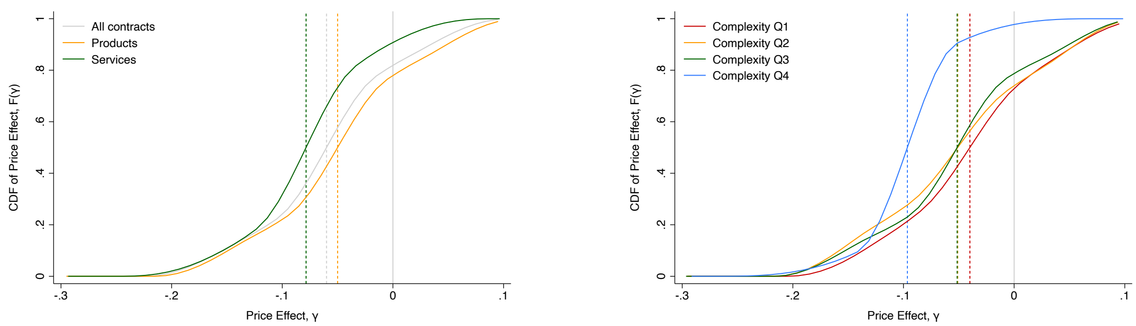

Notes: The panel (a) shows the CDF of price effects separating contracts of goods and services. The panel (b) describes the CDF of price effects by quartile of complexity. This procedure builds upon comparing the empirical densities to a zero-effect counterfactual distribution of publicized distribution under zero effect. The counterfactual distribution is generated from the interpolation of a polynomial of degree 5. The dashed vertical line marks the mean level.

Figure 1.11: Performance effects by sub-group

(a) Reduced-Form Effects on Cost-Overruns by Complexity Quartile

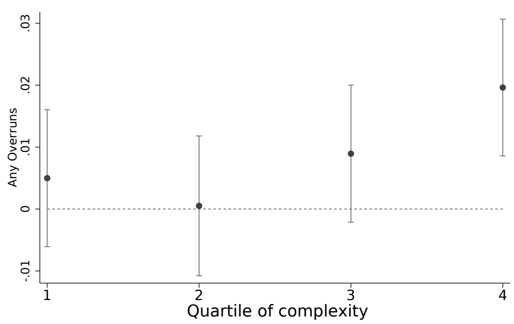

(b) Reduced-Form Effects on Delays by Complexity Quartile

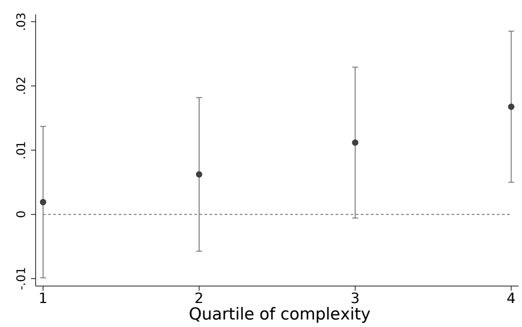

Notes: This figure shows four regression coefficients and their 95% confidence intervals. Each coefficient is a estimate of a RD reduced-form Equation 1.5 per sub-group estimated using (interacted) OLS. The dependent variable of Panel (a) and (b) are indicators for any positive cost-overruns and delays, respectively. The subgroups are determined by the four quartiles of a proxy of contract complexity. The contract complexity proxy is constructed at the product category level and is defined as the average cost overruns for contracts with awards below $20,000 in that category. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019.

Figures

Table 1.1: Summary statistics

	Mean
Contract Characteristics	
Award Amount	20,807
Expected Duration (days)	54.10
Fixed-Price Contract	0.999
Set Aside Award	0.571
Simplified Procedure	0.971
Competition	
Number of Offers	3.542
One Offer	0.239
Contracting Office Characteristics	
Navy	0.378
Army	0.441
Air Force	0.150
Other	0.031
Awarded Firm Characteristics	
Foreign	0.099
Within-State Firm	0.690
Small Business	0.752
Womam Owned Business	0.188
Sample	
No. of Contracts	85,661
No. of Contracting Offices	597
No. of Firms	29,641

Notes: This table presents summary statistics. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. An observation is a contract, defined by aggregating all contract *actions* (initial award, modification, termination, etc.) associated with the same contract ID.

Table 1.2: Top product and service categories

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

Notes: This table presents average annual counts of contracts in the most common product categories. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. An observation is a contract, defined by aggregating all contract *actions* (initial award, modification, termination, etc.) associated with the same contract ID. A 4-digit alphanumeric code (PSC) is observed for each contract. The categories listed are constructed by aggregating PSC codes to two-digits for goods, and to a single digit (letter) for services.

Table 1.3: Estimated Price Effect

Estimate / Sample	All	Goods	Services	Complexity Q1	Complexity Q2	Complexity Q3	Complexity Q4
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean (μ_γ)	0.0595	0.0498	0.0782	0.0397	0.0505	0.0510	0.0962
	(0.0201)	(0.0622)	(0.0596)	(0.0475)	(0.1692)	(0.1908)	(0.0920)
Standard Deviation (σ_γ)	0.0643	0.0670	0.0534	0.0669	0.0739	0.0680	0.0369
	(0.0075)	(0.0084)	(0.0202)	(0.0140)	(0.0760)	(0.0295)	(0.0280)

Notes: This table shows the estimates corresponding to the effect of publicity on contract prices. The estimates result from analyzing the observed contract price density distribution relative to a counterfactual distribution. The observed densities are generated using bins of width $250. The counterfactual distribution stems from a polynomial interpolation of degree 5. The standard deviation is calculated over the non-parametric distribution of γ. The standard errors are calculated through bootstrap. The subgroup analysis is performed independently for each group.

Table 1.4: Reduced-form RDD Estimates and Corrections

Dependent Variable	OLS (1)	CCT (2)	Price Effect Adjustment (3)	Manipulation Bounds (4)	Price Effect + Manip. Bounds (5)
Number of offers	0.3569 (0.0677)	0.5447 (0.1053)	0.3526	[0.2762 , 0.5344]	[0.3073 , 0.4506]
One offer	-0.0191 (0.0064)	-0.0235 (0.0108)	-0.0204	[-0.0272 , 0.0052]	[-0.0248 , -0.0070]
Log distance firm-office	0.1392 (0.0481)	0.1199 (0.0817)	0.1909	[0.0290 , 0.2688]	[0.1304 , 0.2619]
Foreign firm	0.0357 (0.0045)	0.0508 (0.0078)	0.0375	[0.0328 , 0.0520]	[0.0358 , 0.0465]
New firm	0.0175 (0.0075)	0.0185 (0.0126)	0.0247	[0.0023 , 0.0350]	[0.0164 , 0.0344]
Small business	-0.0277 (0.0065)	-0.0295 (0.0110)	-0.0265	[-0.0523 , -0.0195]	[-0.0399 , -0.0220]
Any cost-overrun	0.0135 (0.0045)	0.0246 (0.0077)	0.0144	[0.0103 , 0.0263]	[0.0127 , 0.0216]
Cost-overruns (relative dollars)	0.0095 (0.0058)	0.0161 (0.0100)	0.0127	[0.0053 , 0.0179]	[0.0104 , 0.0174]
Any delay	0.0130 (0.0047)	0.0151 (0.0080)	0.0143	[0.0094 , 0.0270]	[0.0123 , 0.0222]
Delays (days)	2.3262 (2.0388)	4.0361 (3.4935)	2.7491	[1.2278 , 5.3554]	[2.1492 , 4.4660]
Number of modifications	0.0375 (0.0173)	0.0619 (0.0300)	0.0395	[0.0204 , 0.0926]	[0.0300 , 0.0701]

Notes: This table shows Regression Discontinuity Design (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. Coefficients in column (1) use a linear fit above and below the discontinuity. Coefficients in column (2) 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 Appendix A.4. 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 Appendix A.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.

Table 1.5: Expected Ex-Post Adaptations by Product Category

	Goods			Services		
Rank	Name	Average Cost-Overrun	Average Delay	Name	Average Cost-Overrun	Average Delay
Low						
1	Fuels, Lubricants, Oils, Waxes	-0.003	0.009	Transport, Travel, Relocation	0.016	0.029
2	Musical Inst/Phonograph/Home Radio	-0.001	0.016	Construct Of Structures/Facilities	0.026	0.131
3	Valves	-0.000	0.016	Installation Of Equipment	0.027	0.090
High						
1	Chemicals And Chemical Products	0.037	0.062	Operation Of Govt Owned Facility	0.758	0.703
2	Ammunition And Explosives	0.034	0.110	Utilities And Housekeeping	0.343	0.320
3	Office Mach/Text Process/Visib Rec	0.030	0.045	Medical Services	0.270	0.269

Notes: This table presents the top and bottom 3 product categories in terms of average cost-overruns and average delays. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between \$ 5,000 and \$ 20,000, awarded by the Department of Defense in fiscal years 2015 through 2019. We define the fraction of cost-overrun as the final price, including all modifications, minus the award price divided by the award price. The average delay is the final contract duration minus the original contract duration divided by the original duration. These statistics are constructed based on all contracts for the same contract ID. A 4-digit alphanumeric code (PSC) is observed for each contract. The categories listed are constructed by aggregating PSC codes to two-digits for goods and to a single digit (letter) for services.

Chapter 2

Competition under Incomplete Contracts and the Design of Procurement Policies II: Structural Estimates

2.1 Introduction

In the chapter 1 we study the effects of increasing competition through publicity. We find that increasing publicity introduces a trade-off, on one hand it reduces the price at which the contract is awarded, however on the other increases the chances of experiencing cost-overruns and delays. We document that product categories that are more complex are the ones that are more adversely affected by the increase in competition.

This evidence leaves several important research questions that we cannot answer with this reduced-form evidence; in particular, on the buyer side, what are the drivers of the buyer's decision when deciding to restrict competition? How would the buyer behave if the regulation provides more discretion? What would be the optimal design accounting for adverse selection in this setting? To answer these questions, we develop and estimate an equilibrium model of competition for procurement contracts with two general objectives. First, we estimate the underlying firm characteristics, which shape adverse selection in this market. Second, the estimated parameters allow us to study the role of buyer preferences in the promotion of competition, as well as 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.

Firms participate if the cost of preparing the bid is sufficiently low. 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).[1] The award mechanism is a first-price, sealed-bid auction. Fourth, the awarded contractor executes the contract. The quality of execution depends on the existence and magnitude of cost-overruns, which stem from an idiosyncratic shock realized ex-post. The model 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 the buyer's preferences over outcomes, and allows for idiosyncratic preferences for certain vendors that are uncorrelated with contract outcomes.

We estimate our model using data on auction participation, contract prices, and observed cost-overruns. 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 their participation costs are substantially lower; they are more likely to participate, *ceteris paribus*. They are also considerably more prone to experience cost-overruns in the execution stage. Buyers show a preference for lower prices, lower cost-overruns, and incumbent suppliers. We then use our model to estimate the effects 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.

Our results show that imposing regulation to promote bidder participation involves a risk of allowing under-qualified firms to bid. An alternative policy design is to rely on buyers (i.e. each local agency) to decide on whether to publicize each contract. As emphasized by the vast literature on the allocation of authority within organizations (Aghion and Tirole, 1997), 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 in which the buyer decides whether to publicize each contract. We find that delegating this decision to the buyer is welfare-enhancing when the transaction unit is complex: on average, the buyer achieves better outcomes than in regulated settings with either zero or full publicity. However, when the transaction unit is relatively simple, imposing full-publicity rules is convenient as the risks at the execution stage are minor.

We next use our model to identify improvements to the current policy design, which depart from uniform publicity requirements. Policies that regulate competition in most public

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

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 that are tailored to the complexity of the purchase, thus leveraging the benefits of intense competition for simple products, while limiting its adverse effects on complex products. We find that the cost-minimizing level of publicity for commodity-type products is 100%, whereas more complex product categories should require low use of publicity. We find that this reduces average defense procurement costs by 2 percent, or $104 million annually.

This paper contributes to the growing literature that evaluates policies aimed at promoting (or restricting) bidders' participation in procurement settings. This literature emphasizes that expanding the pool of potential bidders may not necessarily translate into lower contract prices if bidders' participation is endogenous, as their equilibrium bidding behavior can become less aggressive (Athey, Levin, and Seira, 2011; Athey, Coey, and Levin, 2013; Li and Zheng, 2009, 2012; Krasnokutskaya and Seim, 2011; Marmer, Shneyerov, and Xu, 2013; Bhattacharya, Roberts, and Sweeting, 2014; Sweeting and Bhattacharya, 2015).[2] We leverage variation in the number of potential bidders that stems from exogenous changes in publicity requirements. We model entry and bidding decisions and find that incumbents are less likely to participate when they anticipate fiercer competition. However, in our setting, the effect of competition from new entrants dominates that of less aggressive bidding by incumbents, reducing the winning bid as a result. The source of variation in the number of potential bidders is closely related to Coviello and Mariniello (2014), who study a similar policy in Italy.

This paper contributes to the growing literature that examines a buyer's role as an agent affecting market outcomes. In particular, this literature considers the fact that buyers' actions can be motivated by objectives other than simple contract price reductions (Bandiera, Prat, and Valletti, 2009; Liebman and Mahoney, 2017; Coviello and Gagliarducci, 2017; Best, Hjort, and Szakonyi, 2017; Decarolis et al., 2020; Carril, 2019; Szucs, 2020). In particular, this paper relates to Kang and Miller (2017), who study buyers' competition promotion for IT contracts in the United States. We depart from existing papers by comparing buyer preferences for ex-ante and ex-post outcomes, with idiosyncratic preferences for specific vendors.

This paper builds upon the context and data discussed in chapter 1. It is organized as follows, in section 2.2, we develop and estimate an equilibrium model of procurement competition. We use the model parameters to study outcomes under counterfactual environments in Section 2.3.

[2]These ideas were initially introduced by Samuelson (1985); Levin and Smith (1994). Li and Zheng (2009) provide an empirical framework highlighting that increasing the number of potential bidders within the independent private value (IPV) setting has ambiguous effects, as the equilibrium behavior interacts two opposite effects: "competition effect" with "entry effect." The former tends to reduce prices, while the latter tends to increase them.

2.2 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 the 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. Section 2.2 describes the equilibrium conditions of the theoretical model. Section 2.2 illustrates the empirical implementation of the theoretical model. In Section 2.2 we discuss key variation in the data to identify model parameters. The results are discussed in Section 2.2.

Model

A buyer offers a single and indivisible contract to N potential contractors. Each potential contractor j, must incur an entry cost $\omega_j > 0$ to learn her private cost to complete the task $c_j \in [\underline{c}, \overline{c}] \subset \mathbb{R}_+$ and, hence, bid for the contract. Both, ω_j and c_j are assumed independent random draws from specific distributions. We model the potential bidders' choices in two stages; first, knowing the number of potential competitors N, each potential bidder decides whether to incur the entry cost. After the entry stage, the $n \leq N$ firms that incurred entry costs learn their costs of completing the job and submit their bids. The awarding mechanism is a first-price sealed-bid auction; the contract is awarded to the bidder that submits the lowest offer. The quality of the contract execution q_j is observed once the contract is finished.

Our analysis considers asymmetry between potential contractors. In particular, there are two types of firms, locals (L) and non-locals (NL). These firms differ in their distributions of entry costs, G^k_ω, the cost of completing the project, F^k_c, and the execution quality F^k_q, where $k \in \{L, NL\}$ and $k(j)$ denotes group affiliation of bidder j. We assume that project and bid preparation costs are private information of each firm and are distributed independently across all firms and identically within group.

Departing from existing literature, our model allows for the set of potential contractors to be chosen endogenously, i.e., the buyer decides whether to publicize the contract taking into account the expected price, the quality of execution and the likelihood that a contractor of each group is awarded the contract. The publicity decision determines the set of potential participants as follows; if the contract solicitation is publicized openly, both locals and non-local contractors learn about the contract solicitation. Conversely, if the contract is not advertised, only the local contractors receive the information.

Equilibrium in the Bidding Stage

We start by characterizing the bidding stage and then use the results to analyze the participation stage. Our analysis focuses on group-symmetric equilibrium where bidders of group k follow the same bidding strategy, $\beta_k(\cdot)$, mapping project cost, c_j, into a bid b_j. Where

c_j is drawn independently from a type-specific continuous distribution $F_c^k(\cdot)$, with density $f_c^k(\cdot)$ and common support $[\underline{c}, \overline{c}] \subset \mathbb{R}_+$. 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 group n_t^k (Li and Zheng, 2009).

Our setting considers two possible scenarios; if the contract solicitation is publicized, then both, *local* and and *non-local* firms could participate. In this case, the expected utility of bidder j with cost realization c_j and group membership $k(j)$ depends on the number of bidders of each group:

$$\mathbb{E}\left[\pi_j(c_j)\right] = (b_j - c_j)\left(\sum_{l=1}^{N^{k(j)}} \rho_l^{k(j)}\left(1 - F_c^{k(j)}\left(\beta_{k(j)}^{-1}\right)\right)^{l-1}\right)\left(\sum_{l'=1}^{N^{-k(j)}} \rho_{l'}^{-k(j)}\left(1 - F_c^{-k(j)}\left(\beta_{-k(j)}^{-1}\right)\right)^{l'}\right)$$

where $\rho_l^{k(j)}$ is the probability that the number of actual bidders is equal to l, and $-k(j)$ denotes the other group of potential contractors. The optimal bidding requires solving a system of differential equations corresponding to the first order conditions for both types of bidders as follows:[3]

$$\begin{aligned}
\frac{1}{b_j - c_j} &= f_c^{k(j)}\left(\beta_{k(j)}^{-1}(b_j)\right)\frac{\partial\beta_{k(j)}^{-1}}{\partial b_j}\left[\frac{\sum_{l=1}^{N^{k(j)}} \rho_l^{k(j)}(l-1)\left(1 - F_c^{k(j)}\left(\beta_{k(j)}^{-1}\right)\right)^{l-2}}{\sum_{l=1}^{N^{k(j)}} \rho_l^{k(j)}\left(1 - F_c^{k(j)}\left(\beta_{k(j)}^{-1}\right)\right)^{l-1}}\right] \\
&+ f_c^{-k(j)}\left(\beta_{-k(j)}^{-1}(b_j)\right)\frac{\partial\beta_{-k(j)}^{-1}}{\partial b_j}\left[\frac{\sum_{l'=1}^{N^{-k(j)}} \rho_{l'}^{-k(j)} l'\left(1 - F_c^{-k(j)}\left(\beta_{-k(j)}^{-1}\right)\right)^{l'-1}}{\sum_{l'=1}^{N^{-k(j)}} \rho_{l'}^{-k(j)}\left(1 - F_c^{-k(j)}\left(\beta_{-k(j)}^{-1}\right)\right)^{l'}}\right] \qquad (2.1)
\end{aligned}$$

If a contract solicitation is not publicized, only *local* firms can bid, i.e., the number of potential non-local contractors is zero. In this case, the bidding problem is symmetric as there is only one group involved. Local suppliers observe contracts' publicity status and hence the number of potential competitors.

Equilibrium of the Entry Stage

Firms compare the *ex-ante* expected profit conditional on entry to their entry cost $\omega_k \overset{iid}{\sim} G_\omega^k$. Firms with entry costs below their expected profit decide to incur the entry fee to learn about their cost of completing the project. The *ex-ante* (expected) profits from participating are given by:

[3] 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. The bidding strategies cannot be solved analytically but require numerical methods. Campo, Perrigne, and Vuong (2003) and Brendstrup and Paarsch (2003) discuss non-parametric identification of cost functions in this setting.

$$\bar{\pi}^k(\varphi^k, \varphi^{-k}) = \int_{\underline{c}}^{\bar{c}} \left(\sum_{n_k-1, n_{-k} \subset N_k-1, N_{-k}} \pi^k(c|n^k-1, n^{-k}) Pr\left(n^k-1, n^{-k}|N^k, N^{-k}\right) \right) dF_c^k(c) \quad (2.2)$$

where φ^k and φ^{-k} are the entry probabilities of each group. Because entry decisions are made simultaneously, the equilibrium condition is characterized by a group-specific entry cost threshold $\bar{\omega}_k$, i.e., firms whose entry cost is below their group's threshold participate.[4] Finally, when the contract is not publicized, only locals could participate, and thus the participation problem becomes symmetric. Thus, for a given contract t, the local's participation threshold differs depending on whether the contract was publicized.

Empirical Model

Based on equilibrium conditions of the general model, we proceed to describe its implementation based on the empirical setting. A contract solicitation t is characterized by (x_t, z_t, u_t, N_t), where x_t, z_t are observable characteristics,[5] and u_t is the unobserved project heterogeneity that captures project attributes that are not included in the data but impact firms' bidding and participation behavior (Krasnokutskaya, 2011). Finally, $N_t = (N_t^L, N_t^{NL})$ denotes the number of potential contractors of each group.[6] The model thus proceeds in four stages depicted in the Figure **??**:

T = 0: *Publicity Decision.* In the first stage the buyer observes (x_t, z_t, u_t, N_t) and decides

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

$$\begin{aligned} \varphi^L &= G_\omega^L\left[\bar{\omega}_L(\varphi^L, \varphi^{NL})\right] \\ \varphi^{NL} &= G_\omega^{NL}\left[\bar{\omega}_{NL}(\varphi^L, \varphi^{NL})\right] \end{aligned}$$

Group-specific equilibrium exist by Brouwer's Fixed Point Theorem. Following Krasnokutskaya and Seim (2011) and Krasnokutskaya (2011), we numerically verified uniqueness of the equilibrium entry probabilities within our estimation routine.

[5] Examples of observable characteristics are the type and complexity of the product required, location, acquiring agency, etc.

[6] Identifying the potential number of bidders is not trivial (Athey, Levin, and Seira, 2011; Krasnokutskaya and Seim, 2011; Mackay, 2018). We combine two methodologies: First, using the procedure described in section 1.3, 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). 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 non-local potential bidders. Finally, we define the number of potential bidders for every buyer-product as the maximum of both approaches. Combining these two methods alleviates potential weaknesses of each of them. The median number of potential local and non-local bidders is six and three, respectively.

whether to publicize the contract to maximize expected utility. Contract publicity status determines the set of potential bidders.

T = 1: *Entry Decision.* In the second stage, each firm that learns about the contract observes (x_t, z_t, u_t, N_t). They draw individual and private realizations of entry cost, and they 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.* The implementation quality is realized once the contract is finished based on a publicly observed quality shock.

Now, we outline specific modelling assumptions of each stage of the model. In Section 2.2 we discuss the estimation strategy to recover the parameters of the underlying distributions from available data.

Publicity Decision. We assume the buyer is risk-neutral and derives utility on contract outcomes: $U_t^D = U(\bar{P}_t^D, \bar{Q}_t^D, \bar{L}_t^D)$, where $D \in \{0,1\}$ denotes contract's publicity status, and $\bar{P}_t^D$, $\bar{Q}_t^D$ and $\bar{L}_t^D$ are the expected awarding price, ex-post implementation quality and likelihood that a local wins, conditional on (x_t, z_t, u_t, N_t). The buyer chooses to publicize a contract t if $U_t^1 \geq U_t^0$

Entry and Bidding Decision. The bidder's j cost for project t is multiplicative: $c_{jt} = \tilde{c}_{jt} u_t$, $\tilde{c}_{jt}$ is a firm-specific cost component that is private information of firm j, while u_t represents a common cost component that is known to all bidders but is unobserved to the researcher (Krasnokutskaya, 2011; Haile and Kitamura, 2019).[7] The distribution of the firm-specific cost component for group-k firms is given by $F_{\tilde{c}}^k(\cdot|x_t)$, and is independent conditional on observable contract characteristics. The unobserved project heterogeneity is given by $u_t \sim H(\cdot)$, is independent from project characteristics and the number of potential bidders.

We assume bidders are risk neutral. Thus, the Bayes-Nash equilibrium bid function for group k is multiplicative: $\beta_{k(i)}(c_{jt}|x_t, u_t, N_t) = u_t \cdot \tilde{\beta}_{k(j)}(c_{jt}|x_t, N_t)$.[8] Each bidder submits a bid of $b_{jt} = \tilde{b}_{jt} u_t$ where $\tilde{b}_{jt} = \tilde{\beta}_{k(i)}(\tilde{c}_{jt}|\cdot)$ represents the bid for bidder j when u_t is one. Therefore, $ln(b_{jt}) = ln(u_t) + ln(\tilde{b}_{jt})$, the log of the unobserved heterogeneity component acts as a additive mean shifter to the conditional distribution of log bids. The contract is awarded using a first-price auction to the bidder that submits the lowest bid.

Finally, we assume the entry costs ω_{jt} are independent (conditional on observed and unobserved project characteristics). The firms' participation behavior is characterized by group-specific thresholds, $\bar{\omega}_t^k(\cdot)$. Thus, the number of actual bidders n_t^k from group $k \in \{L, NL\}$ is distributed according to a binomial distribution with probability of success of

[7]The inclusion of unobserved heterogeneity follows directly Krasnokutskaya (2011). This extension is important to account for the within auction bid correlation that stems from unobserved factors.

[8]Krasnokutskaya (2011) uses deconvolution methods to show that, when the cost function is multiplicative to unobserved heterogeneity, the Bayes-Nash equilibrium bidding strategies can be identified and are also multiplicative.

$\varphi_k(x_t, u_t, z_t, N_t)$ and N_t^k trials: $\varphi^k(x_t, u_t, z_t, N_t) = G_k(\bar{\omega}_{kn}(x_t, N_t), z_t^e))$, i.e., every potential bidder of group k has a probability $\varphi^k(\cdot)$ of entering. Our model considers entry shifters z_t^e, which capture market-level conditions that affect entry decisions.

Contract Execution. The quality of the execution is observed after the contract is performed by the selected contractor. The observed quality *ex-post* is q_{jt}, where q_{jt} corresponds to a draw from group-specific quality distribution $F_q^k(\cdot|x_t, z_t)$. In this section, we focus on one direct measure of performance that is the existence and magnitude of cost overruns. This variable is convenient as it can be directly benchmarked with the contract's dollar value.[9]

Equilibrium is characterized by the buyer choosing contract's publicity status that maximizes her expected utility; informed potential contractors entering if expected profits exceed entry costs and bidding optimally in the market mechanism. Finally, the quality of the implementation is revealed once the contract concludes.

Estimation Approach

In estimation, we make functional form assumption to characterize the distributions of interests.[10] We assume that the log of individual bids $log(\tilde{b}_{jt})$ are distributed normal with mean $\mathbb{E}[\tilde{b}_t^k|x_t, N_t] = [x_t, N_t]'\alpha^k$, and variance: $\mathbb{V}[\tilde{b}_t^k|x_t, N_t] = (exp(x_t'\nu^k))^2$. We further assume $ln(u_t)$ is distributed normal with mean zero and variance σ_u^2.[11] In equilibrium, the entry decision is characterized by a type-specific probability, $\varphi_t^k(x_t, z_t^e, N_t)$, which depends on a type-specific entry-cost distribution (Krasnokutskaya and Seim, 2011; Athey et al., 2011, 2013). We assume $\varphi_t^k(x_t, z_t^e, N_t,) = \Phi([x_t, z_t^e, N_t]'\tau^k)$, where $\Phi(\cdot)$ denotes the cumulative distribution of the standard normal distribution, and z_t^e are entry cost shifters. The number of participating bidders $n_t^k(\cdot)$ is distributed binomial with N_t^k independent draws with probability of success $\varphi_t^k(\cdot)$.[12]

The quality of the implementation is captured by the existence and magnitude of cost overruns. Given that most contracts stay right on budget, the observed distribution of cost overruns is censored at zero. We assume that $ln(q_t^k)$ is the latent distribution, while we only observe $Q_t^k = max\{0, ln(q_t^k)\}$, where $ln(q_t^k)$ distributes normal with mean $\mathbb{E}[q_t^k|x_t, z_t^q] = [x_t, z_t^q]'\gamma^k$, and variance $\mathbb{V}[q_t^k|x_t] = (exp\{x_t'\xi^k\})^2$. Where z_t^q denotes cost-overrun mean shifters.[13]

---

[9] This approach abstracts away from other (context-specific) execution costs. For example, Lewis and Bajari (2011) study the welfare gains associated with reducing delays on high-way construction.

[10] Our parametric assumptions are linked to related literature. Overall, our results are not sensitive to adding additional covariates or variations to the functional form. Our data provide enough variation for identifying these distributions independent from the functional form.

[11] These parametric assumptions follow existing 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.

[12] Related papers either assume parametric distributions on the entry costs, which, paired with the expected utility of entering, map into well-defined group-specific entry probabilities (Krasnokutskaya and Seim, 2011; Mackay, 2018), or make functional form assumptions on the entry probabilities, which, combined with expected utilities, allow recovering entry costs (Athey et al., 2011, 2013). We follow the latter approach.

[13] Given the structure of the model, we are assuming that the cost-overruns capture excess of the cost, which the

Finally, an important feature of this model is that we aim to recover buyer's preferences combining estimated parameters with observed publicity decisions. Buyer's utility is decomposed into observed and unobserved parts; the observed part is assumed to be a linear combination of expected outcomes, in particular, we define $\left(\tilde{\bar{P}}_t, \tilde{\bar{Q}}_t, \tilde{\bar{L}}_t\right)$ as the change in the expected outcome under publicity and no publicity, leaving expected outcome without publicity as the omitted category. We further assume the buyer has idiosyncratic preferences distributed standard normal. This way, $Pr(D_t = 1|\cdot) = \Phi\left(\beta^P \tilde{\bar{P}}_t + \beta^Q \tilde{\bar{Q}}_t + \beta^L \tilde{\bar{L}}_t + x_t'\zeta\right)$, where β^P, β^Q, and β^L capture the relevant taste parameters for the price, quality, and idiosyncratic preference for awarding local vendors.[14] We control for observable characteristics as well as an indicator of whether the expected price without publicity is above the regulation threshold. The latter enters in the utility function as a discrete shift of publicity's utility.[15]

Sample. The data used to estimate the model is the same as those used in previous sections with the exception that to classify local and non-local vendors, we require buyer-product combinations that appear at least four times in the full database between 2013-2019, with at least one, but not all, appearances in FBO. This restriction rules out products that are purchased less often. Table B.2.1 compares the descriptive statistics of relevant variables between this selected sample and the full sample used in Section 1.3. Overall, given that the sample selection involves the buyer contracting the same product multiple times, the selected sample includes contracts for categories that are, on average, less durable, i.e., over-represent services. This sample selection does not affect the main results presented in Section 1.3. Finally, and consistent with the rest of the analysis, we estimate the model using contracts around the regulation threshold, i.e., between 10 and 40 thousand dollars.

Estimation. Our empirical model yields predictions about equilibrium conditions for suppliers' participation, bidding, and ex-post execution with and without publicity. Also, we characterize the buyer's publicity decision. Our estimation strategy proceeds using simulated method of moments (Mcfadden, 1989; Pakes and Pollard, 1989). That is, we choose a vector of parameters θ to generate simulation-based moments that closely resemble key moments from the data. Using parametrized primitives discussed previously, we simulate four empirical objects: participation decisions, bidding strategies, quality delivered and publicity decisions and create a set of moments conditions to be matched with data.

Our simulation procedure starts with a set of size T of data inputs (x_t, z_t, N_t), then from every observation we generate S random draws of u_t. Finally, our setting contemplates that the buyer decides based on expectations, these expectations are formed conditional on (x_t, z_t, N_t) and u_t, integrating over Monte Carlo simulated distributions of price, quality,

vendor can entirely transmit to the buyer. Thus, there's no strategic behavior on behalf of the contractor nor the buyer ex-post. The fact that the observed contractor's overruns are drawn from group-specific distributions speaks about differences in cost prolixity and knowledge about buyer's context. We assume these differences are exogenous.

[14] Our estimation does not restrict the set of values for these parameters; however, in general, we should expect that buyers dislike paying higher prices or experiencing overruns, so we expect β^P and β^Q to be negative.

[15] Intuitively, the larger the utility shift generated by the regulation, the higher will be the observed jump of publicity adoption at the threshold.

and the likelihood of a local winning. This method, although computationally involved, is useful to circumvent integrating over potentially non-linear functions, and provides enough flexibility to match theoretical moments functions that cannot be evaluated directly.

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

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \left(m_n - m_s(\theta)\right)' W_n \left(m_n - m_s(\theta)\right) \tag{2.3}$$

Where W_n is the weighting matrix, which is chosen using the standard two-step approach; the quasi-optimal weight matrix W_n is derived in the first stage, and the parameters are estimated in the second stage (Gourieroux, Monfort, and Renault, 1993). The vector of parameters corresponds to: $\theta = (\alpha^k, \nu^k, \tau^k, \gamma^k, \xi^k, \vec{\beta}, \zeta, \vec{\sigma})$.

We use three sets of target moments. The first set of moments are a vector of first and second-order moments of the relevant variables as well as it's interaction with it's relevant covariates. The relevant outcome variables are the auction price, the number of bidders, local winner, the magnitude of cost-overruns, an indicator of any cost-overrun, and publicity choices. The second set of moments consist of means of the same outcome variables conditional on partitions of the domain of contract prices. These moments capture the relation between these outcome variables over the domain of prices and are estimated separately for goods and services. Finally, the third set of moments are a vector of normalized frequencies on the relevant window of contract prices. Stacking together these three vectors, we obtain the vector m_n of 357 moments that seek to match with the model. We use the stochastic optimization algorithm *Differential Evolution* (Storn and Price, 1997) to perform the objective minimization.[16] The details of the estimation procedure are discussed in the Appendix B.3.

Identification

The model identification involves identifying primitive distributions of the two types of bidders (locals and non-locals). In our setting, however, the contract's publicity status determines the composition of participating bidders; if a contract is not publicized, only locals could participate, while if it's advertised, both types can compete for the contract. Data of unpublicized contracts inform distributions of locals, and thus, non-local distributions are pinned down from publicized contracts.

We assume that contract covariates, entry and overruns shifters, as well as the number of potential bidders (x_t, z_t, N_t) are exogenous. We leverage (conditional) independence between u_t and c_{jt} and that $\mathbb{E}[u_t|x_t, z_t, N_t] = \mathbb{E}[u_t] = 1$ to identify the expected proportional bids, cost

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

shocks distributions and relative profits (Krasnokutskaya, 2011). In our setting, exogenous variation on expected competition due to variation in N or entry shifters, impact the private cost component of the winning bid, but not the common cost u_t. The latter, combined with the first-order condition (equation 2.1), pins down the proportional cost distribution (Guerre, Perrigne, and Vuong, 2000; Campo et al., 2003), and thus, firms' expected profit. The latter, linked with observed entry decisions, allows us to pin down the distribution of entry costs. The distribution of unobserved heterogeneity is obtained after the distributions of individual cost shocks are identified.[17]

The contractor's execution (cost overrun) distribution stems from an individual quality shock. We assume the quality shocks are independent (conditional on observables). Thus, the observed distribution of cost-overruns informs the underlying distribution that generates quality shocks for each bidder group.

Identifying type-specific distributions from the variation on contract publicity choices can only be achieved if the latter is exogenous. In the spirit of RDD discussed before, we leverage the discrete nature of the publicity threshold to obtain quasi-experimental variation on publicity adoption and thus get exogenous variation in the publicity decision to identify type-specific distributions separately. The buyer's taste parameters for price, overruns, and local contractors are identified based on the exogenous nature of contract and market characteristics (x_t, z_t, N_t). In particular, variation on entry and on the number of potential bidders of each type determine the publicity effects 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 well-defined (commodity-type) product, there would be no difference in performance ex-post, which shuts down that factor in the decision.

While our estimation approach hinges on parametric assumptions of primitives' distributions, these distributions can be identified nonparametrically based on distributional assumptions discussed above. Nonparametric identification, along with the reduced-form results and robustness checks, suggests that the estimated distributions' features are not driven by functional form.

Estimation Results

Estimation of the model proceeds in two steps. In the first step, we estimate the model's parameters of entry, bidding, performance and publicity selection. In the second step, we combine the estimates with model equilibrium conditions to recover the distribution of production and entry costs for locals and non-locals. These estimates are inputs to the policy counterfactuals in section 2.3.

[17] Our model considers endogenous entry, and we leverage additional variation in entry shifters, to pin private and common cost distributions. However, separate identification of private and common costs distributions can be achieved without modeling entry as long as N is exogenous (i.e., no selection on unobservables). See Mackay (2018).

We specify the mean of log bids as a linear function of the product's degree of complexity , an indicator of service, and the numbers of potential bidders of each group. We interact the complexity with a dummy for non-local bidders to allow the effects of most of these covariates to differ by bidder's group. The variance of log-bids depends on whether the group is a good or service. The probability of entry depends on the same covariates, and we add a dummy to indicate the solicitation was required at the end of the fiscal year.

The cost-overruns' shock is distributed log-normal; its mean depends on the product's degree of complexity, an indicator of service, interactions with non-local dummies, and an indicator if the contract duration exceeds the mean.[18] Finally, the buyer's decision to publicize a contract depends on the expected differences (in logs) of price and overruns, as well as the predicted likelihood that a local wins. In addition to these variables, we include agency fixed effects and a dummy if the contract's expected value without publicity is over $ 25,000.

Estimates

Table 2.2 displays our estimates with the corresponding standard errors.[19] To facilitate the interpretation of coefficients, table 2.3 shows the marginal effects for the set of coefficients associated with the bidders and the buyers.

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 1.3, non-local contractors are 72 p.p. more likely to participate than locals. Furthermore, this is consistent with the fact that bidders' reduce the probability of entering if they observe more potential non-local competitors; one additional non-local reduces the probability of participating by 7.4 p.p.

Second, bids from non-locals are slightly lower than locals; they bid 4 p.p. lower prices. 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 auction. The standard deviation of (log) unobserved heterogeneity is 27 times larger than the bids' standard deviation when $log(u_t) = 0$.

Third, the quality shock depends on the transaction product; the mean of log-quality shocks is substantially higher for services and complex product categories. In line with reduced form results, the difference between locals and non-locals in execution quality is substantial; non-locals have a mean shock that is 23 p.p. higher. Interestingly, the difference between these two groups is relatively stable over different degrees of product complexity.

Finally, as discussed in previous sections, contract publicity allows non-local bidders to participate in auctions, which leads to different contract outcomes. Panel C of Table 2.3 shows that the buyers choose to publicize 2.4 p.p. more if they anticipate that publicity leads to a 10% reduction in awarding price. A 10% increase in cost overruns reduces the probability

[18] The mean duration is calculated using contracts under $20,000 to remove the influence of the threshold.

[19] Although the model is estimated altogether, tables 2.2 and 2.3 present estimates in different columns to facilitate visual interpretation. The procedure to estimate standard errors is discussed in Appendix B.3.

of advertising by 1 p.p. Buyers have a preference for local vendors; if they anticipate a 10 p.p. reduction in the likelihood of a local winning, buyers reduce the probability of publicity by 2.3 p.p. Finally, predicting that the price without advertising exceeds $25,000 increases the likelihood of publicity by 32 p.p. The latter is in line with the increase in probability estimated in Section 1.3. These estimates depart from the standard assumption that the buyer *only* aims to minimize price and provides valuable inputs to evaluate policy counterfactuals.

Model Fit. Overall, the model closely replicates the key empirical patterns in the estimation sample. We examine model fit using the estimated parameters to simulate equilibrium outcomes and compare simulated to observed outcomes. Our simulations, and the ones discussed later, build upon the estimating dataset by drawing simulations of the unobserved heterogeneity and the quality, entry, and bidding-cost shocks. These draws, combined with estimated parameters, allow us to simulate market-level equilibrium. Figure B.1.2 compares the the distribution of model-simulated outcome variables with actual data. The simulated data replicates closely publicity choices, actual bidders, and the share of contracts assigned to locals. Panels (e) and (f) separate cost-overruns by products and services. We find that, 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.

Recovering the Cost Distribution

We recover the distribution of project costs leveraging the methodology introduced by Guerre, Perrigne, and Vuong (2000), and Campo, Perrigne, and Vuong (2003). This method combines the first-order conditions (equation 2.1) —subject to boundary conditions— with estimated equilibrium bids to estimate the inverse bid function. Since, in our setting, the actual number of bidders is unknown, the first-order condition depends on the probabilities of different combinations of local and non-local bidders. These probabilities are formed from simulations based on model's participation parameters. Finally, strict monotonicity between the bid and the inverse bid functions enables us to to obtain an estimate of project costs from estimated distribution of bids.

Figure 2.1a shows the probability density function of log costs ($log(\tilde{c}_{jt})$) of both groups. Local bidders have slighly higher costs than non-locals. Figure 2.1b displays the mean $log(\tilde{b}_j(\tilde{c}))$ as a function of the log cost. As expected, markups decrease with higher the cost draws.

Recovering Entry Costs

We recover the group-specific entry-costs using the equilibrium conditions for optimal entry behavior discussed in section 2.2. A potential bidder compares the draw from the entry-cost distribution G^k_ω with the expected utility of entering, i.e., $\varphi^k(x_t, z_t, N_t) = G^k_\omega\left(\bar{\pi}^k(x_t, N_t), z_t\right)$.

Our estimated cost distributions $F_{\tilde{c}}^{k}(c)$ allow us to estimate the (*ex-ante*) predicted utility of participating (equation 2.2) and benchmark it to observed entry behavior (Athey, Levin, and Seira, 2011).

Our estimated cost distributions imply that the distribution of entry costs differ substantially between the two groups. On the one hand, roughly 60% of non-locals face zero entry cost (i.e., they enter with probability one), and 90% enter if the entry cost is less than 0.1 log units. On the other hand, local contractors face substantially higher entry costs. Indeed, with a 60% chance, they would *not* enter for any of the values included in the estimated range of existing expected utilities. The estimated entry-cost asymmetry shapes the composition of actual bidders and, subsequently, the winning bid conditions.

Figure B.1.3 in the Appendix shows how the composition of actual competitors (and the identity of the winner) depends on the number of potential non-locals. It shows how the number of actual bidders decreases as the number of potential non-locals increases. This is consistent with the fact that reductions on predicted utility due to increased competition discourage locals' participation, making it substantially more likely that a non-local wins.

Implications Increasing Competition through Publicity

Having estimated the primitives of the model as a function of observable characteristics, we replicate the policy evaluation discussed in section 1.3 and, evaluate contract's outcomes with and without publicity, not only for contracts around the threshold, but throughout the values included in our sample.

Figure 2.3 displays the variation on contract outcomes around the regulation threshold as a function of the expected price without publicity and simulate the results in a setting without publicity thresholds. These results are in line with our reduced-form analysis discussed previously. Publicizing contract solicitations allows the participation of non-local bidders, which are more prone to experience overruns and have substantially lower participation costs discouraging locals' participation. Thus, enhancing contract participation through publicity reduces prices ex-ante due to increased competition; however, it increases prices *ex-post*. We propose the following definition of the final price that takes into account both effects:

$$p_{D,t}^{F} = p_{D,t}^{I}(1 + q_{D,t})$$

where $D \in \{0,1\}$ denotes contract t's publicity status, $p_{D,t}^{I}$ is the log awarding price, and $q_{D,t}$ is the realized share of of cost-overruns *ex-post*. Thus, $p_{D,t}^{F}$ denotes contract t's log final price.

Figure 2.4, compares the consequences of publicizing contracts at different levels of complexity. Auctions for complex contracts have a higher variance of bid functions and face lower participation levels. The former increases the support of possible price reductions by added bidders. Adding one bidder to an auction with lower participation levels has more effect than adding one on an auction that already has many competitors. Thus, these two market features contribute to more extensive effects of added competition on award prices.

Our results show substantial asymmetry between local and non-local vendors when executing contracts. Non-locals experience cost-overruns considerably more often than locals, i.e., contract publicity leads to higher cost-overruns ex-post. Therefore, the net of these opposing effects depends on the degree of complexity of the transaction; for relatively complex contracts, the increase in cost-overrun exceeds *ex-ante* price reductions, i.e., more competition leads to higher final contract prices. Conversely, there's little to no rise in cost-overruns associated with publicity for simple contracts. Thus publicizing contracts leads to reductions in final contract costs.

These findings align with our reduced form results and reinforce the idea introduced by seminal papers on incomplete contracts; when the transaction is involved, and the number of possible contingencies ex-post is higher, assuring proper ex-post performance becomes more important than reducing prices *ex-ante* (Williamson, 1976; Bajari and Tadelis, 2001; Bajari et al., 2014; Bolotnyy and Vasserman, 2019).

2.3 Counterfactual Analysis

We leverage estimated model parameters to evaluate the implications of counterfactual policies on award prices and ex-post cost overruns. Overall, there are two approaches to improving outcomes in a principal-agent setting; (1) delegating the publicity decision to the buyer relying on the buyer's knowledge of the local market and context, or (2) impose rules that restrict the agent's span of possible actions. Our counterfactual exercises aim to benchmark these two approaches. In particular, we evaluate what would be the set of actions that a buyer would take in a deregulated setting or as a result of alternative regulation designs.

The Strategic Value of Delegating Competition Promotion to the Buyer

What are the implications of allowing the buyer to choose contracts' publicity status relative to tighter publicity rules that prescind the buyer's decision? This trade-off pertains to the more general problem of the delegation of authority within organizations (Aghion and Tirole, 1997), and it's a frequent theme of debate in procurement policy discussion (Kelman, 1990).

Conceptually, the publicity requirement design works as a discontinuous jump in the cost of not publicizing; below the threshold, the cost of not advertising is zero, whereas above the threshold is positive and involves filing additional paperwork. Therefore, below the threshold, buyers publicize their desired publicity with full discretion, while above the threshold, they are forced to advertise more than desired (due to the added cost of not publicizing). Thus, using the estimated model parameters, we back out the buyer's hypothetical decisions in a full discretion setting and evaluate its implications relative to the current policy design.

Figure 2.5 evaluates the implications of current publicity requirements for different levels of contract complexity. Panel (a) shows how the fraction of publicized contracts increases at the threshold relative to a counterfactual scenario without the requirement. Panel (b) illustrates how the price-complexity relation shifts as a result of increased publicity.

The effectiveness of allowing buyers to decide whether or not to publicize contracts depends on the product's complexity. Buyers with full discretion obtain lower prices than what they would get with no publicity over all the complexity spectrum. However, if contract complexity is lower than 0.18, a full publicity setting achieves lower average prices. Hence, providing discretion effectively reduces costs if the transaction unit is complex, whereas forcing publicity is more effective if the unit of transaction is relatively simple.

The current regulation combines these two scenarios; it provides discretion below the threshold and nudges higher levels of publicity above the threshold. As a result, relative to a regulation-free setting, the current regulation involves higher levels of advertising, which is effective when the purchase is simple; however, it backfires when the transaction is complex. In section 2.3, we discuss the effects of alternative policy tools to improve contract outcomes.

The Role of Buyer's Preferences

In our setting, the buyer decides the degree of contract competition motivated by interests that are not necessarily the same as that of the organization. The agency problem hinges on the degree of misalignment between the buyer's (agent) and the organization (principal) objectives. Thus, the issue disappears if their objectives coincide since the buyer, by maximizing her utility, would be, too, maximizing the organization's welfare. An extensive theoretical literature studies the design of compensation mechanisms to "align" agent's objectives (Laffont and Tirole, 1993). This literature often builds upon rationality and completeness of contract menus to improve outcomes.

We study the extent to which contract outcomes depend on specific preference parameters. To do so, we depart from the estimated parameters, varying the degree of "alignment" on the buyer's preferences. We set two benchmark situations: First, the buyer has "Cost-Oriented Preferences", i.e., puts equal weight on price reductions ex-ante and ex-post and has no idiosyncratic preference for local contractors. Second, the buyer has "Local-Oriented Preferences", i.e., preferences oriented to favor local contractors with no emphasis on costs. The specific preference parameters under each scenario are described in Table 2.4. It is worth mentioning that these two benchmark preference parameters are based on the estimated coefficients –"shutting down" specific taste parameters. Therefore, they should be seen as reference points of policies oriented to affect buyers' motives.

Figure 2.6 shows (log) final price effects depending on the level of complexity of the purchase. Buyers with "Cost-Oriented Preferences" decide to publicize to reduce costs, so, and perhaps not surprisingly, relative to other preference schemes, they generate savings all across the spectrum of product complexity. "Local-Oriented" agents seek to benefit local contractors; they publicize infrequently, and, as a result, the outcome prices are similar to the situation without any publicity.

Existing literature on rules vs. discretion (Aghion and Tirole, 1997; Carril, 2019; Bosio, Djankov, Glaeser, and Shleifer, 2020) emphasizes that regulation can be an effective antidote to waste and abuse whenever these are pervasive. Still, it can backfire if agents are virtuous in exercising discretion. Our findings contribute to the existing literature by highlighting that this trade-off depends, too, on the level of contract complexity: Publicity regulation can be detrimental even when agents are misaligned, as favoring local vendors has positive spill-over on cost-overruns. Moreover, we find that there is room for improving outcomes when buyers are aligned through strict publicity requirements when the transaction unit is simple.

Complexity-Based Publicity Requirements

We now follow a more tactical approach and take the estimated buyer's preferences as given and vary the regulation design because the effects of publicizing depend on the level of contract complexity. The proposed design contemplates identifying the "right level" of publicity requirements depending on the degree of contract complexity. In this exercise, we refer to publicity requirements as the minimum fraction of contracts that buyers must publicize. 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 price at each complexity level. As a result, the proposed regulation imposes publicity requirements that are specific to each product category depending on the complexity level.[20]

Figure 2.7 summarizes this procedure: Panel (a) shows the price-complexity relation at different publicity requirements. Panel (b) illustrates the publicity requirement that minimizes price at different complexity levels. Panel (c) shows the price-complexity relation that would stem from a "tailored" publicity requirements. Note that the latter (brown-dashed line in panel (c)) corresponds to the lower contour of final prices at different requirements (panel (a)). These tailored publicity requirements alter the span of the buyer's actions. In particular, when the unit of purchase is simple, it removes the buyer's discretion entirely to leverage the benefits of enhanced competition; however, it provides more choice when contracts are more complex to attenuate the negative consequences on contract implementation ex-post.

[20] The proposed design involves a higher degree of regulation sophistication as it requires fixed fractions of publicized contracts per complexity level. We believe that can be implemented smoothly given the set of rules included in the Federal Acquisition Regulation. As a reference, the current version of chapter 5 (Publicity Requirements) 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 simple product, whose optimal level of publicity requirement is 100%, 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.

Comparing Policy Counterfactuals

Table 2.5 brings together the visual evidence provided Figures 2.5, 2.6 and 2.7, and compares the overall mean price effect under each of these scenarios. The current policy design that introduces publicity requirements at \$ 25,000, reduces, on average, the final price by 1.6% relative to a no-publicity scenario. If the publicity choice were delegated to the buyer, this reduction would be, on average, 1.3%. A uniform full-publicity rule would reduce contract costs by 3%. However, the latter could be improved if the publicity requirements were tailored to the purchase's degree of complexity. The latter policy reduces prices by 3.6%. The 2% cost difference between the current and the complexity-based design corresponds to \$104 million — competitively awarded— defense contracts, annually.[21]

Finally, we benchmark the consequences of these policy designs with the hypothetical situation in which the buyer has cost-minimizing preferences. We find that the cost-minimizing publicity requirement achieves better outcomes than Cost-Oriented Buyers. From a policy standpoint, this is significant as arguable modest improvements to the regulation design could mitigate most of the concerns associated with misaligned buyers and achieve, on average, better outcomes than any compensation mechanism that aims to align buyers' objectives.

2.4 Conclusion

In this chapter, we develop and estimate an equilibrium model of competition of procurement contracts, with two general objectives. This model allows us to estimate the underlying firms' characteristics that shape adverse selection in this market and buyers' objectives when promoting competition through advertising. Furthermore, our estimates allow us to evaluate the implications of alternative policy designs. We focus on two relevant counterfactual policies. The first we assess the implications of delegating competition promotion to the buyer. The second aims to propose a welfare-enhancing regulation design that accounts for the vast heterogeneity in the degree of complexity of procurement transactions.

Our counterfactual analysis shows that delegating competition promotion to the buyer is only welfare-enhancing when the transaction unit is complex: on average, the buyer achieves better outcomes than in regulated settings with either zero or full publicity. However, when the transaction unit is relatively simple, imposing full-publicity rules is preferred as the risks at the execution stage are minor. Moreover, 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 will significantly reduce procurement costs. Notably, while our analysis is carried out using data from the Department of Defense, we believe that the general conclusions apply broadly to private and public organizations' transactions.

[21] This amount is calculated using defense contracts competitively awarded in 2018 with values between \$10,000 and \$150,000 (Simplified Acquisition Threshold).

Figures

Figure 2.1: Estimated Cost Distributions

(a) Cost Distribution (b) Bidding Function

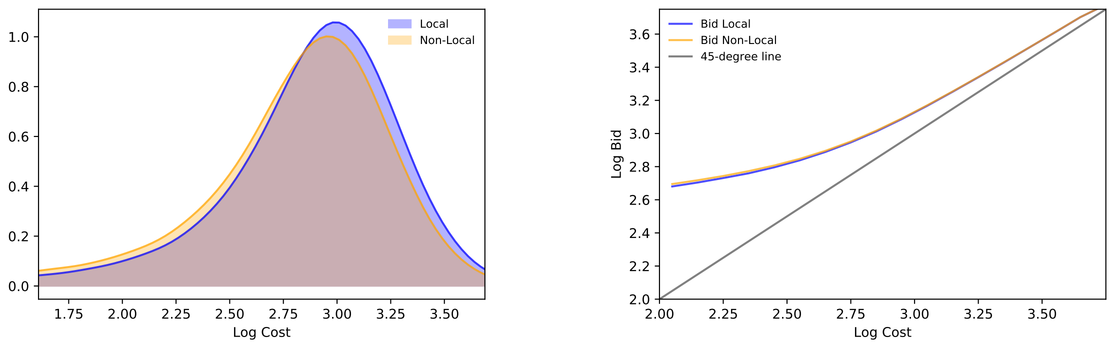

Notes: The panel (a) shows the distribution of log costs for locals and non-locals. The panel (b) displays the bidding function. Both plots are estimated with average covariates and $log(u) = 0$. The plotted distribution of log costs is smoothed using a kernel.

Figure 2.2: Estimated Entry Cost Distributions

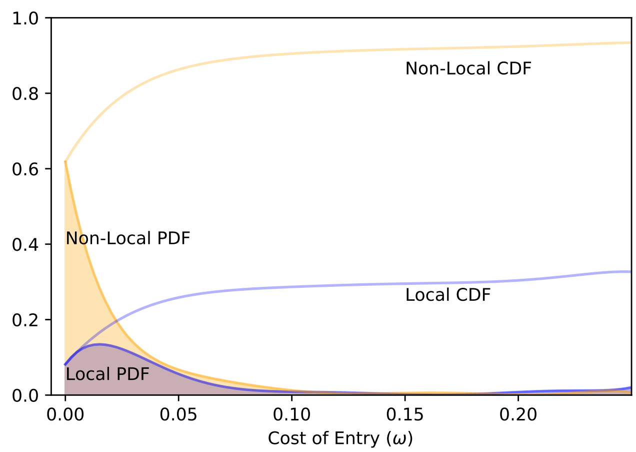

Notes: This figure shows the estimated entry cost probability and cumulative density distributions with average covariates and $log(u) = 0$. The CDFs comparing the entry probability to different levels of predicted (log) utilities. The PDFs capture the differences of the CDFs within bin widths of size 0.025. The raw distributions are smoothed with a 7-degree polynomial fit. By nature, zero, in this case, is the minimum expected utility.

Figure 2.3: Policy Evaluation Using Model Parameters

(a) Number of Bidders

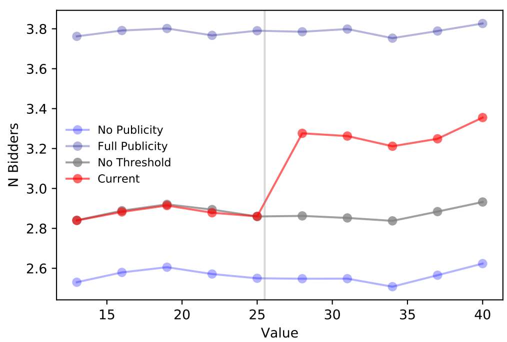

(b) Local Wins

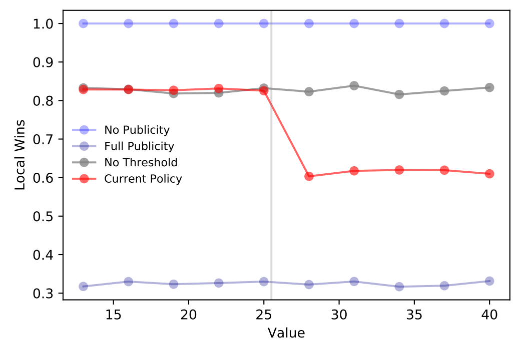

(c) Log Prices

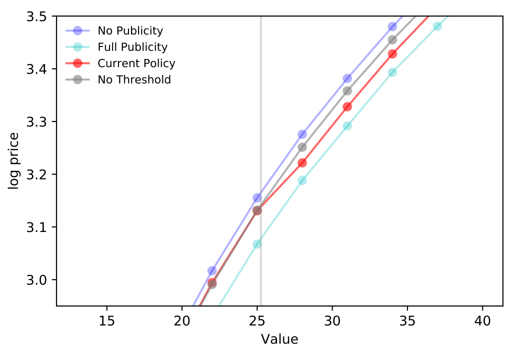

(d) Overruns

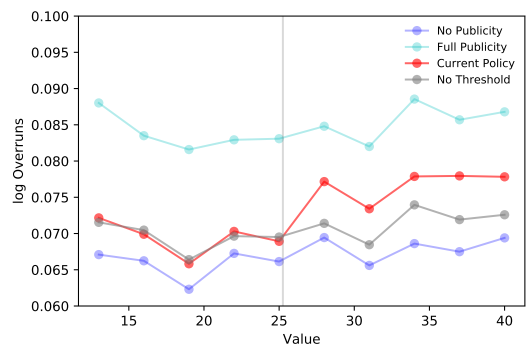

Notes: These figures show how different outcomes vary around the threshold. Panel (a) shows the number of bidders, panel (b) the probability of awarding a local, panel (c) the log of price, and panel (d) the log overruns. The x-axis of every graph is the value of the contract without publicity, which us the relevant dimension affected by the regulation threshold. As a reference, in every graph, we benchmark the current policy design (red line) with a situation of full (and zero) publicity. Also, we include the simulated event of no threshold as a counterfactual scenario.

Figure 2.4: Final Price by Product Complexity

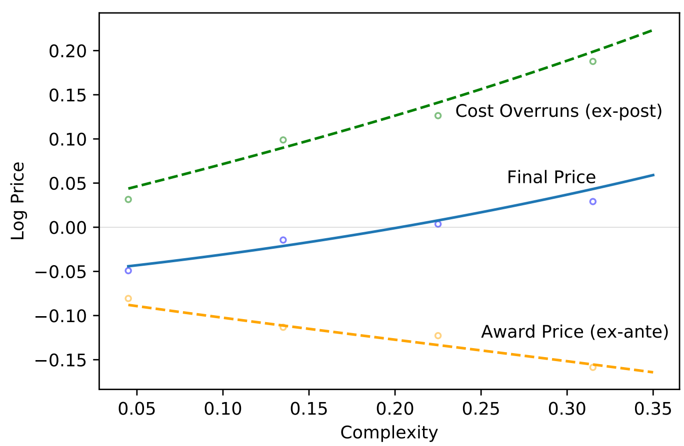

Notes: This figure shows the effect of publicity on (log) of award price, on (log) of cost overruns, and (log) of the final price for different degrees of product complexity. The omitted category is the variable without publicity. The circles represent the mean effect per complexity bin. Each line corresponds to a flexible polynomial fit. The degree of complexity is defined as the log of the product's average overruns, and it is calculated on all contracts for the same product category below $20,000.

Figure 2.5: Final Price by Product Complexity

(a) Publicity Adoption

(b) Log Final Price Effect

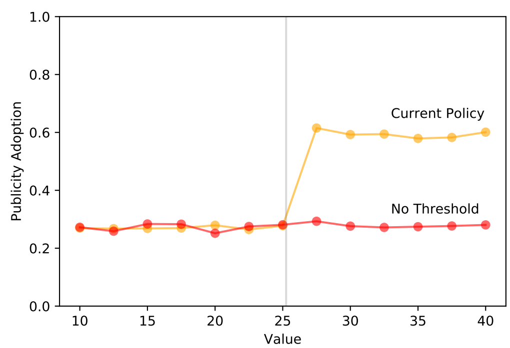

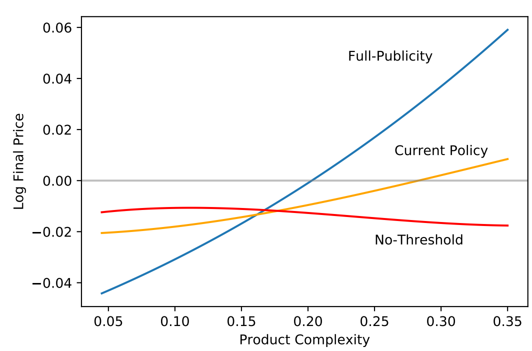

Notes: The panel (a) shows the share of publicized contracts depending on the value of the contracts without publicity. The orange line displays the current policy with a regulation threshold at \$25,000. The red line describes the share of adoption in the absence of the regulation threshold. The panel (b) shows the effect on the log of the final price ($p^F_{m,t} - p^F_{0,t}$) for different degrees of product complexity. The blue line is the effect if all contracts are publicized, the orange line is the effect of the current policy (with a threshold at 25,000), the red line shows the effect in the absence of regulation threshold. The omitted category is the log of the final price in the absence of publicity. Each line was constructed using flexible polynomial fits. The degree of complexity is defined as the log of the product's average overruns, and it is calculated on all contracts for the same product category that are below the regulation threshold (\$25,000).

Figure 2.6: Final Price by Product Complexity

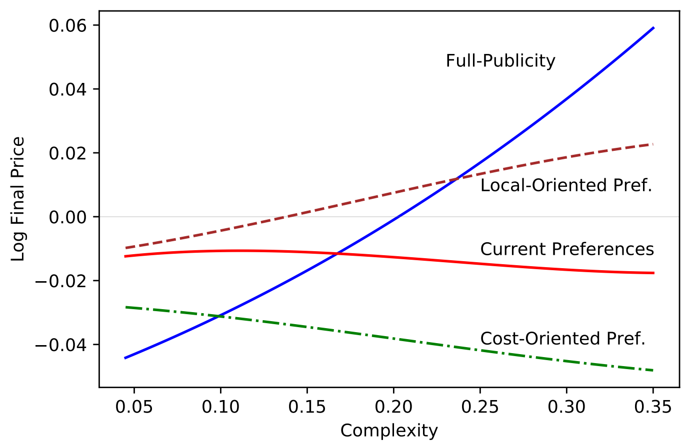

Notes: This figure shows the effect of publicity on the log of the final price for different preference parameters. The blue lines represent the price effect if all contracts are publicized. The red line is the price effect under current preference parameters. The green and brown dashed lines represent the effects of cost-oriented and local-oriented buyers' preferences, respectively. All lines are flexible polynomial fits. The degree of complexity is defined as the log of the product's average overruns, and it is calculated on all contracts for the same product category that are below the regulation threshold ($25,000).

Figure 2.7: Final Price by Product Complexity

(a) Publicity Adoption

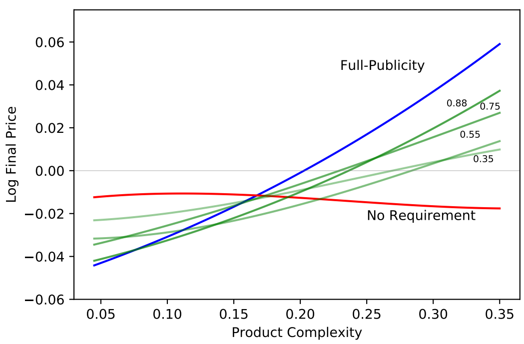

(b) Log Final Price Effect

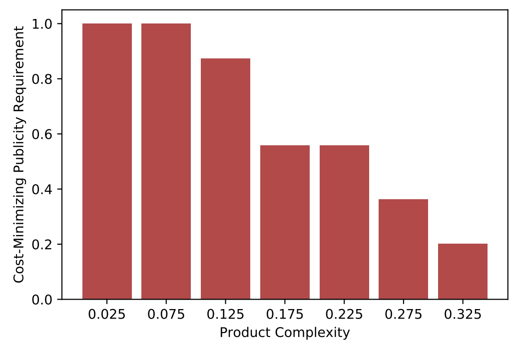

(c) Log Final Price Effect

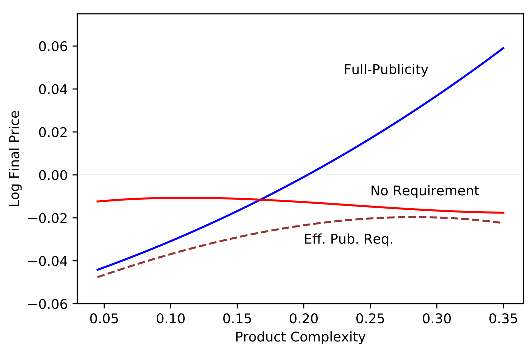

Notes: The panel (a) shows the log of price effect with full publicity (blue line), without regulation requirements (red line), and for different levels of publicity requirements (green lines). The panel (b) describes the publicity requirement that yields the minimum price for varying levels of complexity. Overall, the publicity requirements that minimize costs decrease with product complexity. The panel (c) shows the price effect when imposing efficient publicity requirements at different levels of complexity. As a result, the price effect of this policy (brown-dashed line) corresponds to the lower contour of the panel (a). The omitted category is the log of the final price in the absence of publicity. Each line was constructed using flexible polynomial fits. The degree of complexity is defined as the log of the product's average overruns, and it is calculated on all contracts for the same product category that are below the regulation threshold ($25,000).

Tables

Table 2.1: Summary Statitistics Local vs. Non-Locals

	Local	Non-Local	Diff
log Distance	3.471	4.554	-1.083
Located in the Same State	0.695	0.501	0.194
Overruns (relative)	1.078	1.236	-0.158
Delays (relative)	1.130	1.275	-0.145
Number of Modifications	0.548	0.880	-0.332

Notes: This table presents summary statistics for distance and execution variables for contracts performed by locals and non-locals. The sample includes contracts between 10,000 and 40,000 dollars, and buyer-product combinations that appeared at least four times between 2013 and 2019. The need for observing multiple buyer-product observations stem from the way we categorize these contractors. The variables "Overruns" and "Delays" are the ratio of the final relative to the original cost and duration, respectively. The difference between these groups is significant at 1% for all variables.

Table 2.2: Estimated Parameters of Entry, Bidding and Execution

		Entry (Probit)		Bid Distribution (Log Normal)		Execution (Log Normal)	
	$\bar{x}$	Coeff	S.E.	Coeff	S.E.	Coeff	S.E.
Panel A: Coefficients							
Constant		0.0048	(0.00057)	3.0517	(0.00025)	-1.7041	(0.00038)
Service	0.375	-0.0598	(0.00014)	0.0000	(0.00030)	0.0751	(0.00026)
Degree of Complexity	0.089	-0.7367	(0.00071)	-0.0005	(0.00037)	1.6284	(0.00066)
Non-Local		2.1651	(0.00040)	-0.0402	(0.00037)	0.2354	(0.00069)
Non-Local×Complexity		0.0299	(0.00026)	-0.0308	(0.00063)	0.0868	(0.00082)
Last Month	0.249	-0.8826	(0.00077)				
Exp. Duration ¿ Median	0.5					0.1273	(0.00025)
N^L	6.078	0.0002	(0.00010)	-0.0023	(0.00027)		
N^{NL}	3.339	-0.1876	(0.00028)	0.0099	(0.00024)		
Panel B: Standard Deviation							
Contant				-2.6132	(0.00081)	-0.4102	(0.00048)
Service				0.0996	(0.00031)	1.1854	(0.00048)
S.D. Unob. Het. (σ_u)				2.1683	(0.00087)		
Panel C: Buyer Preferences							
		Publicity Choice (Probit)					
Constant		-0.2584	(0.00059)				
Exp. Price (β^P)		-0.6361	(0.00043)				
Exp. Cost-Overruns (β^Q)		-0.2457	(0.00058)				
Exp. Local Winning (β^L)		0.5879	(0.00069)				
Above \$25K		0.8542	(0.00037)				
Number of Obs.	24,135						

Notes: The table displays the coefficients and corresponding standard errors. Panel A describes the coefficients corresponding to the entry choice mean of (log) bids and the mean of (log) quality shocks. Panel B displays information estimates for the standard deviation of (log) bids, unobserved heterogeneity, and (log) quality. Panel C shows the coefficients associated with the publicity choice by the buyer. Agency and year fixed effects are omitted in this table. These coefficients are estimated altogether using Simulated Method of Moments (SMM). Log-bids and the log of the unobserved project heterogeneity are assumed to be normally distributed. The entry and publicity choices distribute Probit. The standard deviation of log-bids and log-quality shocks are estimated as $\sigma = exp(b_0 + b_1 \mathbb{1}(Service))$, where $\mathbb{1}(Service)$ indicates a contract for service

Table 2.3: Marginal Effects Model Estimates

	$\bar{x}$	Δx	Entry $\Delta\varphi/\Delta x$	Bidding $\Delta\tilde{b}/\Delta x$	Execution $\Delta q/\Delta x$	Execution $\Delta(q>0)/\Delta x$
Panel A: Marginal Effects						
Service	0.38	1	-0.024	0.000	0.075	0.013
Degree of Complexity	0.09	0.1	-0.028	0.000	0.163	0.030
Non-Local		1	0.721	-0.040	0.235	0.039
Non-Local×Complexity		0.1	0.001	-0.003	0.009	0.001
Last Month	0.25	1	-0.333			
Exp. Duration ¿ Median	0.50	1			0.127	0.021
N^L	6.08	1	0.000	-0.002		
N^{NL}	3.34	1	-0.074	0.010		
Panel B: Standard Deviation						
Estimated ($\hat{\sigma}$)				0.076	1.035	
Unob. Het. (σ_u)				2.168		
			Publicity			
Panel C: Marginal Effects		Δx	$\Delta Pub/\Delta x$			
Exp. Price		0.1	-0.0243			
Exp. Cost-Overruns		0.1	-0.0095			
Exp. Local Winning		0.1	0.0228			
Above $25K		1	0.3263			
Number of Observations	24,135					

Notes: Panel A shows the marginal effects of main on different dependent variables related to bidders' actions. The Marginal Effects are computed at the mean of each covariate described in the second column. The third column shows the change used to estimate the effect. The dependent variables are the probability of entry, the bid level, and the quality shock in terms of levels and the probability of over zero. Panel B shows the empirical models' estimated standard deviation and the estimated standard deviation of the unobserved heterogeneity component. Panel C displays the marginal effects of expected price, cost-overruns, local winning, and being above $25K on the probability of publicizing the contract solicitation in *FBO.gov*. These coefficients are jointly estimated using Simulated Method of Moments (SMM).

Table 2.4: Buyer's Preferences

	Estimated	Benchmarks	
	Preference Parameters	¡4-¿Cost-Oriented ¡4-¿Preference	Local-Oriented Preference
β^P	-0.636	¡4-¿-0.636	0
β^Q	-0.245	¡4-¿-0.636	0
β^L	0.588	¡4-¿0	0.588
Mean Pub.	0.274	¡4-¿0.408	0.263

Notes: The first column shows the estimated preference parameters for the price ex-ante, overruns ex-post and awarding local contractors. The second column shows the preference parameters associated with a buyer with cost-oriented preferences, i.e., no idiosyncratic preference for locals. The third column shows the preference parameters for a buyer that is fully oriented to locals, i.e., does not have a preference for price or quality but only for favoring local contractors. The last row, 'Mean Publicity,' describes the average adoption of publicity under each of these types of preferences.

Table 2.5: Effects of Counterfactual Scenarios

	Mean	95 Perc C. I.
Current Policy	-0.0164	[-0.0195 , -0.0133]
Delegate the Decision to the Buyer		
Current Preferences	-0.0136	[-0.0161 , -0.0111]
Buyer with Price-Oriented Pref.	-0.0304	[-0.0335 , -0.0271]
Buyer with Local-Oriented Pref.	-0.0068	[-0.0092 , -0.0043]
Alternative Regulation Designs		
Full Publicity	-0.0306	[-0.0357 , -0.0256]
Complexity-Based Publicity Req.	-0.0367	[-0.0412 , -0.0323]

Notes: The table reports the mean equilibrium effects to log final prices under different scenarios. In the first row, describes the estimated effect under the current policy, with threshold that increases publicity at $ 25,000. In the second row, the mean effects without publicity requirements. The third and fourth row show the mean price reduction if under cost-oriented and local-oriented buyers, respectively. The fifth row shows the mean price reduction under full publicity. The sixth row shows the average price reduction implementing efficient publicity requirements depending on the level of complexity of the product. These effects are constructed relative to the situation without publicity. Confidence intervals are constructed via bootstrap.

Chapter 3

Slippery Fish: Enforcing Regulation when Agents Learn and Adapt

3.1 Introduction

Correcting market failures and improving economic efficiency often require curbing undesirable behaviors of market agents who act to maximize their private benefits. Examples span actions that affect the natural environment, such as deforestation, pollution, or resource exploitation (Stavins, 2011; Duflo, Greenstone, Pande, and Ryan, 2013, 2018; Hansman, Hjort, and León, 2018); Actions that affect community health such as open defecation or drunk driving (Banerjee, Duflo, Chattopadhyay, Keniston, and Singh, 2017); Or actions that undermine government performance such as corruption or tax evasion (Carrillo, Pomeranz, and Singhal, 2017). Enforcing regulations is the most direct strategy to deter such behaviors. Enforcement not only requires strong state capacity, but also sophisticated policing to track agents' reactions to audits, so that policies are robust enough to deter cheating even when agents try to 'game' the new system.

Targeted agents adapt to new rules, finding loopholes that allow them to continue maximizing private benefits at the expense of others.[1] In many instances, it is therefore insufficient to evaluate the effectiveness of enforcement activities based on their immediate, short-run effects (Fudenberg and Levine, 2020). A more sophisticated evaluation will need to track the (sometimes unanticipated) strategies that targeted agents may deploy to circumvent the regulation as they adjust to the new regime.

We develop a model of enforcement paired with an experimental design and data col-

[1] For example, Carrillo, Pomeranz, and Singhal (2017) show that when the Ecuadorian tax authority improves the quality of their information on firm revenues, the firms react by raising their estimates of costs in line with the revised revenue estimates, to keep total tax payments unchanged. Blattman, Green, Ortega, and Tobón (2017) shows that intensive policing pushes crime around the corner, with null impacts on overall violent crimes. Health officials adapt to undermine a monitoring scheme to punish delinquent nurses in Banerjee, Duflo, and Glennerster (2008), making an initially-effective program completely ineffective in 18 months.

lection strategy that delineate how agents learn about the patterns of, and loopholes in, enforcement. We highlight adaptation along two different margins: (i) the agent learning audit patterns and schedules over time, and (ii) the agent devising defensive strategies to avoid paying fines even when he is audited. Our Bayesian model of learning also yields predictions on the specific design of enforcement strategies that will be more robust to the agent's subversive adaptation efforts. We test these predictions by conducting a large-scale randomized controlled trial (RCT) in which government monitors penalize vendors that sell illegal fish in Chile, while we surreptitiously monitor vendors' reactions to that enforcement by deploying "mystery shoppers" in fish markets. We cross-randomized the frequency of enforcement visits and the (un)predictability of schedules to test theoretical predictions about the optimal design of audit policy. Not surprisingly, we see that fish vendors find it more difficult to adapt when monitoring visits are unpredictable. The theory and experiments also deliver a counter-intuitive result: Auditing is more effective when it is conducted at *lower* frequency. Even if enforcement is very cheap to conduct, the auditor will sometimes do better by holding back some effort.

This experiment takes place at scale in a consequential setting, targeting a problem that has large welfare consequences. The government of Chile has instituted a ban on fishing and sales of critically endangered Pacific hake fish (*merluza*) during September each year, when the fish reproduces. Catching hake during that period is especially ecologically destructive. We randomize the fish markets where the government sends monitors to levy penalties on vendors illegally selling fish. We also cross-randomize a consumer information campaign designed to educate consumers about the environmental risk associated with over-fishing of hake and discourage consumption during the September ban period. This serves as a useful benchmark because less direct strategies such as information campaigns designed to change social norms around the undesirable behavior[2], or marketing that appeals to people's sense of fairness (Hainmueller, Hiscox, and Sequeira, 2015), or encouraging third-party reporting (Naritomi, 2018), may be more cost-effective in settings where it is difficult to enforce rules. Our information campaign could even complement the audit strategy: If vendors react to the enforcement by hiding their illegal hake sales, then informed consumers may be an important second line of defense. Our 2x2 experimental design can test for such complementarities.

Since we are tracking illegal activities, we measure outcomes using "mystery shoppers" to improve credibility of the data. We sent trained surveyors who look like typical shoppers to each market to pose as buyers and (try to) purchase fish during the ban. We link the daily reports from mystery shoppers to the enforcement logbook recorded by government inspectors to test our model's specific predictions on the nature of learning and adaptation in response to variable patterns of enforcement visits that different vendors experience.

We also conducted consumer surveys to gather data on changes in demand for hake and other substitutes, and consumer knowledge about the hake ban. We mapped all spatial and

[2]For example, Chetty, Mobarak, and Singhal (2014) partners with the Bangladesh tax authorities in an attempt to change social norms to encourage firms to pay taxes (as opposed to enforcing tax laws directly), and Guiteras, Levinsohn, and Mobarak (2015) attempt to change social norms around toilet use (as opposed to directly banning the dangerous practice of open defecation).

market relationships between vendors and fishermen to study spill-overs across markets.[3] Finally, we surveyed the fishermen who supply to these markets to explore whether interventions implemented "downstream" (at the point of sale from vendors to consumers) traveled "upstream" the supply chain of fish. It is ultimately the fishermen who make the ecologically sensitive decisions in the seas. Our sample covers all major markets where the majority of hake is caught, which allows us to report on equilibrium outcomes, such as changes in fishermen activities, or availability and price of hake substitutes. This produces a more comprehensive evaluation of the full range of effects up and down the supply chain.

Our analysis proceeds in three steps. First, we conduct a program evaluation of the government's audit and information campaigns. These interventions lower, but do not eliminate, illegal hake sales. Second, we specify a model of learning and test its predictions, to develop a more precise understanding of how regulated agents learn about the audit system, adapt, and develop defensive strategies. Our mystery shoppers systematically record the new practices vendors introduce to circumvent enforcement. Many do not display the hake openly during the ban, but are willing to sell our mystery shoppers illegal fish that is hidden from plain view. They also start keeping the hake on ice, and claim that the fish on display was caught in August when it was still legal to do so. These reactions attenuate the effects of enforcement on the *true* availability of illegal hake in markets.

Third, we introduce experimental variations in the design of the audit system to test which strategies are more robust to such subversive adaptation. Audits on a predictable schedule become less and less effective over time, as vendors learn monitoring schedules and shift sales away from targeted days and markets. We also tried increasing monitoring frequency to better contain temporal and spatial spillovers to other days of the week or other nearby markets, but this strategy backfires. Increased frequency evidently allowed fish vendors to learn about the flaws in the system more quickly and react with greater hiding and freezing of illegal fish.

Our findings shed light on a larger theoretical literature in Law and Economics on adaptation and subversive reactions to regulations (Glaeser and Shleifer, 2003; Becker, 1968; Eeckhout, Persico, and Todd, 2010; Lazear, 2006). Also related is the literature on *gaming* incentive schemes where agents adapt to undermine the intent of the regulator (Ederer, Holden, and Meyer, 2018; Oyer, 1998; Gravelle, Sutton, and Ma, 2010). That literature suggests that introducing unpredictability and opacity to incentives can mitigate gaming by the agent and improve payoffs for the regulator. Both our model and Okat (2016) predicts that random and less frequent enforcement hinders or delays agents' learning about the weaknesses of the auditing process. Our results call into question any enforcement mechanism that economic theory deems "most efficient" without grappling with the (potentially unanticipated) behavioral responses by regulated agents. We grapple with the real-world complexities of implementing a large government enforcement program at scale, and contribute to the empirical literature on the effects of monitoring and penalties (Boning, Guyton, Hodge, Slemrod, Troiano, et al., 2018; Shimshack and Ward, 2005; Gray and Shimshack, 2011; Hansen, 2015;

[3] We use the gendered term "fishermen" because every single fisher we interviewed in Chile was a man.

Pomeranz, 2015; Johnson, Levine, and Toffel, 2019). The most closely related paper to ours is Banerjee et al. (2017)'s policing intervention to curb drunk driving in India, where they randomized fixed vs rotational checkpoints (akin to our unpredictable monitoring) as well as the frequency of monitoring. They estimate a model of driver learning, and infer evidence of strategic responses.

Beyond regulation and enforcement, we show that an easier-to-implement consumer information campaign is almost as effective in curbing the illegal activity as direct monitoring.[4]

We generate evidence on the real world challenges to implementing an auditing scheme in one specific sector, but the sector and policy we study are globally relevant and important.[5] FAO (2014) estimates that 31.4% of the world's fish stocks were over-exploited to biologically unsustainable levels in 2013, up from 10% in 1974. Costello, Ovando, Hilborn, Gaines, Deschenes, and Lester (2012) reports that over-exploitation is worse in small-scale fisheries like the one we study, and such fisheries represent the majority of the global catch. Illegal fishing accounts for US$10-23 billion worth of fish each year. Fishing bans of the type we study in Chile are in effect in many countries around the world, including China, Fiji, India, Ghana, Bangladesh, Peru and Myanmar. Some of these other policies are extremely similar in structure to the Chile hake ban, such as a 22-day ban on selling Hilsa fish in Bangladesh during the fish's reproduction period, and a 60-day ban on silverfish in Peru.[6]

The paper is organized as follows. We describe the context and experimental design in section 3.2. Section 3.3 develops the theory of learning. Section 3.4 describes data collection. Section 3.5 presents the empirical strategy and results, section 3.6 documents spillovers and market equilibrium effects, and section 3.8 concludes the paper.

3.2 Background and Experimental Design

Context

With around 4,000 miles of coastline, Chile is one of the top ten fish producers in the world (FAO, 2014). However, as in many other low and middle-income countries, the marine ecosystems have been threatened by over-fishing. The Chilean government has passed various regulations to protect threatened species over the last 20 years, including restrictive fishing quotas and fishing ban periods. However, the fish population has continued to decrease, with 72% of species rated as overexploited or collapsed by 2015 (Subpesca, 2015).

[4] Like our consumer information campaign, many other papers have evaluated indirect strategies in pursuit of social goals, in environments where enforcement is expensive or difficult (Johnson, 2016; Jin and Leslie, 2003; Reinikka and Svensson, 2005; Alm, Jackson, and McKee, 2009; Shimeles, Gurara, and Woldeyes, 2017; Kollmuss and Agyeman, 2002).

[5] FAO (2007) emphasizes that "90 percent of the 38 million people recorded globally as fishers are classified as *small-scale*, and an additional more than 100 million people are estimated to be involved in the small-scale post-harvest sector."

[6] See http://www.newagebd.net/article/52220/22-day-ban-on-hilsa-fishing-from-oct-7 and https://elcomercio.pe/economia/peru/produce-establece-veda-nacional-pejerrey-60-dias-noticia-543012

The majority of people carrying out fishing activities are small-scale and artisanal fishermen. Small-scale fishermen contribute almost 40% of the national fishing volume, and up to 75% of the hake fish market. Artisanal fishermen are organized in fishing villages called *Caletas*. Around 76% of the caletas are located in rural areas along the extended Pacific coast, and they are highly spatially dispersed (Subpesca, 2013). Their geographic dispersion, informality, and the small-scale of operations of each individual fisherman make it difficult for the government to monitor their activities. The absence of alternative income-generating activities for these fishermen has also make it difficult to change the norms regarding "appropriate behavior" in this industry. Furthermore, poor small-scale fishermen do not readily accept government-imposed restrictions, and they have organized and unionized to create political opposition to government policies that restrict fishing.

The Pacific Hake is the fish low and middle-income Chileans consume most, and also one of the most important sources of protein for this population. The domestic hake market is served entirely by the domestic supply. Imports and exports of hake are quite uncommon. In an effort to protect the hake population, the Chilean National Marine Authority (*Sernapesca*) and the central government have enacted various policies including restrictive fishing quotas and a one-month ban on fishing and selling hake during the fish's September reproduction cycle. Due to difficulties in enforcing the ban, the hake population is now critically threatened, and has shrunk to 18% of its long-term sustainable level (Subpesca, 2015).

Supply Chain of Illegal Fish

Caletas: Coastal Villages where Artisanal Fishermen Bring in their Catch

Most of the illegal hake fish is captured by small-scale rural fishermen operating out of hundreds of *caletas* dotting the coastline. Each *caleta* contains between 10 and 100 fishing boats. Boats are about 20 and 30 feet in length, and operated by two to three fishermen (see Figure C.1.1). The fishermen operating out of each *caleta* are organized as a union to internally distribute the fishing rights allocated to that *caleta*. In practice, each fisherman captures illegal, undeclared fish beyond the allocated quota. WWF (2017) estimates that the amount of hake fished by small-scale artisanal fishermen are between 3.8 and 4.5 times the legal quota. As a result, the artisanal sector is responsible for 75% of the hake fish supplied in the market, even though they hold only 40% of the "official" hake quotas.

Fishermen go fishing using artisanal boats and nets at night and sell fish after sunrise. They are able to target specific fish types by varying the location and depth at which the nets are dropped. The fish is sold directly at the docks to three types of buyers: (1) fish vendors who buy the fish to sell them in local markets, (2) intermediaries who supply fish to vendors located in places further from the coast, and (3) households who live close to the *caleta* and buy the fish for their own consumption. There is very little use of ice and refrigeration at this point in the supply chain. The fish that vendors sell in local markets is typically fresh, and captured the night before. Table C.3.3 in the Appendix describes *caleta* characteristics.

Ferias: Outdoor Markets where Hake is Sold

The majority of hake-fish sales to final consumers occur in *ferias*, which are outdoor urban markets organized by municipalities. Each vendor pays a fee every six months to rent a selling spot in the market. In addition to fish, *ferias* sometimes contain stalls offering fruit and vegetables, clothes and other products.

Ferias are typically navigable only by foot, and each feria serves a limited geographic area of surrounding neighborhoods. To cover more neighborhoods, the vendors rotate between different *ferias* in a pre-set pattern - typically setting up in the same location twice a week. For example, they may sell at a first feria every Sunday and Wednesday, at a second feria every Tuesday and Friday, and at a third feria every Thursday and Saturday. The group of vendors who move together across neighborhoods is called a *circuit*. A semi-annual fee paid by the vendor to the municipality covers her inclusion in the entire circuit, so the same group of vendors typically rotate across neighborhoods all together. Vendors are not allowed to sell in public places other than *ferias*.

Each municipality typically organizes one circuit of vendors. Large municipalities may have more than one circuit. In such cases, the municipality area is divided in such a way that there is no geographic overlap between circuits. Figures C.1.2 and C.1.3 in the Appendix provide visual examples of ferias and circuits. Table C.3.1 describes observable characteristics of fish stalls in ferias.

Experimental Design

This study was implemented in close collaboration with the Chilean National Fish Service (Sernapesca), who has the ultimate regulatory authority over fishing activities in the country. Our implementing partner's goal from this project was to limit hake fishing, sales and consumption during the September ban. It is practically and politically very difficult for them to directly regulate fishermen, because their activities occur out in the water at night, and because the fishermen operating out of the geographically dispersed *caletas* are politically organized. *Sernapesca* therefore expressed an interest in exploring options to better regulate the fish sales at *ferias* where hake is most commonly sold.

Sample

We conduct our experiment in the five central regions of Chile, which is home to 74% of the Chilean population. The caletas located along the coastal villages and cities scattered across these five regions account for 98% of all hake fish harvested in Chile. We conduct our experiment in all ferias in these regions except for the city of Santiago.[7]

An important benefit of conducting the experiments at such a large and comprehensive scale is that it allows us to track any displacement of illegal hake sales towards control markets, because all potential markets (including ones where the interventions were not applied) are in our database. This allows us to trace the market-level equilibrium effects of our interventions. We collected data on the universe of circuits in our sample area, and from every fish vendor operating in those circuits. We mapped all ferias served by each caleta where the fish are caught. The unique long and thin geographic shape of Chile means that ferias are generally located very close to the caletas from where they source fish (22 miles away on average). This made it relatively easy to connect vendors to the fishermen they source from, and trace how the effects of our interventions are transmitted along the supply chain for hake fish.

There are 280 ferias (fish markets) operating in the 70 municipalities in our sample, and these ferias are organized into 106 separate *circuits*. In order to identify and map all existing ferias and circuits, we combined administrative data from multiple sources (Ministry of Economics and Sernapesca) along with information gathered from phone conversations with staff in every municipality. We then used Google Maps to define the consumer "catchment area" for each feria. We identify the neighborhoods which are likely served by each feria, considering the walking distance and road accessibility from the neighborhood to the feria, as well as the residential versus commercial/industrial characteristics of the neighborhoods. The location of the ferias and their organization as circuits were important for the design of our enforcement intervention. The definitions of the residential neighborhoods and their connections to each feria were important for the design of our consumer information campaign.

Interventions

This study experimentally evaluates the effects of two interventions that aimed to reduce illegal sales of hake during the September ban period. **Enforcement** targeted the **supply** of hake by monitoring vendors and enforcing penalties on those found to be selling illegal hake. The **Information Campaign** was designed to sensitize consumers about this environmental problem and discourage hake consumption during the ban, in order to lower the **demand** for hake.

Design of Enforcement Intervention

The supply-side enforcement intervention deployed government officials from *Sernapesca* to periodically visit ferias where fresh hake is usually sold, and levy fines if vendors are caught

[7] Santiago is unique in that there is one big centralized fish market called *Terminal Pesquero Metropolitano* (TPM) where vendors buy from intermediaries to re-sell at neighborhood ferias. TPM is already well-monitored by *Sernapesca*, and our interventions therefore did not need to be implemented there.

illegally selling hake during the September 2015 ban period. A enforcement visit consisted of two Sernapesca officials visiting all fish stalls in a market. The officials were instructed to follow the usual Sernapesca protocols to search for illegal fish at each stall.[8] Our conversations with vendors prior to September 2015 suggested that they were already well aware of the hake ban. The most important change in 2015 compared to earlier years was that the enforcement activities were applied more consistently and regularly. As a part of this randomized controlled trial, Sernapesca agreed to conduct this monitoring at specific locations and according to schedules defined by the research team. Sernapesca shared the details about their monitoring activities with the research team. The punishment for illegal sales is a US $200 fine plus confiscation of the illegal fish. $200 is equivalent to two weeks of earnings for the average feria vendor, so this represents a significant threat.

We anticipated that fish vendors would react to the enforcement activity by devising new defensive strategies that would help them avoid paying fines. We introduced random variations in enforcement policy design to investigate mechanisms that may be robust to agents' efforts to circumvent policy:

1. *Predictability:* We randomly varied the ease of predictability of the enforcement. In some areas, Sernapesca monitors followed a consistent schedule (e.g. M,W at 9am) while in other areas, they were asked to follow a less predictable schedule defined by the research team. The research team randomly varied the day in which the visit is deployed in any given week, keeping constant the total number. The latter is a more expensive enforcement strategy because it requires having monitors on-call for longer windows. This strategy was therefore practically more difficult for Sernapesca to implement.

2. *Frequency:* We randomly varied audit frequency at the circuit level, so that some groups of vendors only received one visit per week, while others were visited multiple times at the various locations in the city where they set up on different days of the week. Increasing the frequency of monitoring visits is more expensive, but it may limit vendors' ability to relocate illegal hake sales spatially and inter-temporally during the week. On the other hand, it may also accelerate vendors' learning about audit patterns, and deploy effective defensive actions more quickly.

Enforcement activities were randomized at the circuit-level, covering all 106 market-circuits. This randomization was stratified to ensure balance with respect to a few important spatial and market characteristics: Whether the circuit (a) was located in a coastal municipality, (b) was the only one operating in its municipality, and (c) served geographically isolated communities.

[8]The enforcement protocol used in September 2015 was 'business as usual', with no additional instructions to the inspectors. The study design was negotiated at a higher level, and most of the inspectors did not know about the existence (or aim) of the study. They merely followed instructions on where and when to visit markets.

Design of Information Campaign

The demand-side intervention was a marketing campaign designed to inform consumers about the September ban on hake sales. Sernapesca distributed letters, flyers and hanging posters in the residential neighborhoods randomly assigned to this intervention. The message contained in the flyers and posters was simple: "In September, Respect the Hake Ban." The letter, signed by the Director of Sernapesca, included three paragraphs. The first paragraph informed readers about the hake ban every September. The second noted the decline in the hake population to a critical level as a result of over-exploitation, and the third encouraged consumers to not consume hake this month. Appendix C.1 shows samples of flyers and the letter. In previous years, Sernapesca used a smaller budget to place informational ads in newspapers and highway billboards. Information was distributed directly to consumers at a household level for the first time in 2015.

Using our mapping exercise described in section 3.2, and combining it with the location of major roads and crossings, we define boundaries of neighborhoods and divide the municipality up such that the population-size of neighborhoods would be roughly equal. We conducted this intervention in the 48 most populated municipalities and identified 270 distinct neighborhoods in those municipalities. Figure C.1.7 provides example maps. The randomization procedure was as follows:

1. First, 18 of the 48 most populated municipalities were assigned to a high saturation information treatment, 17 to a low saturation information treatment, and the remaining 13 municipalities did not receive the letters, flyers or posters. "High saturation" was defined to be a case where two-thirds of the neighborhood in the feria's catchment area would receive the letters, flyers, and posters. In the low saturation treatment area, only one-third of the neighborhoods received those mailings. We randomly varied the proportion of neighborhoods receiving the treatment to examine whether there are larger changes in norms regarding the acceptability of inappropriate or socially harmful behavior when households observe that many of their neighbors simultaneously receive the same information about the illegality of hake consumption.

2. Second, specific neighborhoods within each high or low saturation information treatment area were randomly chosen to receive the treatment.

3. Third, we randomly selected around 200 addresses in each of 102 neighborhoods, and mailed out letters to each of those 20,400 addresses. 200 letters cover roughly 15% of all potential addresses in a representative neighborhood. Based on information from the postal service, we subsequently learnt that at least 13,000 letters were correctly delivered.[9] 80,000 flyers were distributed by trained field personnel to people walking

[9]Although 13,000 were explicitly tracked, it is likely that around 16,500 were actually delivered, because the postal service did not receive any delivery failure notice in those cases. We inferred and constructed addresses using Google maps, and many of those addresses did not actually exist. That was a leading cause of delivery failure.

in the streets, and directly to households within the 102 treated neighborhoods. 3,000 posters were placed around treated neighborhoods where they would be publicly visible, such as at bus stations, community centers, and street intersections.

Cross-Randomized Experimental Design

The enforcement treatment and the information campaign were cross randomized in a 2x2 experimental design so that we could study potential complementarities between the two approaches. The first panel of the table 3.1 lists the number of circuits assigned to each of the four treatment cells.

The majority of markets were assigned to Enforcement because that column contains additional sub-treatments in which we conduct experiments on variation in enforcement policy design. Those variations in predictability and frequency of enforcement visits were cross-randomized so that we have sufficient statistical power to study the effect of each variation, one at a time. The scond panel of the table 3.1 shows the number of circuits assigned to each sub-treatment cell. To study the effects of predictability of enforcement, we will compare the 39 circuits where Sernapesca monitored on a predictable schedule against the 44 circuits where they monitored on an unpredictable schedule. Similarly, to study the effects of audit frequency, we will compare the 34 circuits assigned to high-frequency against the 49 circuits assigned to low-frequency.[10]

Tests of the information campaign saturation effect (i.e. proportion of neighborhoods around markets that are simultaneously sent letters and flyers), will compare the 30 circuits randomly assigned to a low-saturation campaign (where a third of neighborhoods received letters and flyers), against the other 26 to a high-saturation campaign. We are able to control for other dimensions of random assignment whenever we focus on the effects of one particular dimension. Each of our treatments could have spillover effects on control markets, and we discuss those issues in section 3.6.

3.3 Model of Enforcement

We formalize the decision-making process of a vendor who chooses whether to sell hake illegally, in order to develop empirical predictions we test with daily data from markets. In the process, we develop insights on the nature of learning and adaptation.

Setup

A risk-neutral vendor chooses whether to sell illegal hake in each period $t \in \mathbb{N}$. Selling hake has a fiduciary benefit of $v > 0$. Government inspectors periodically visit the vendor, and if hake is detected, levies a monetary fine $\Omega > v$. The vendor's selling decision depends

[10] The probability of assignment to low-frequency enforcement and to un-predictable schedules was a little higher compared to other cells. In our analysis, we will control for these differences.

on her *perceived* probability of receiving an enforcement visit that day, as well as on the likelihood of being fined if visited. The vendor can adopt (costly) defensive actions to reduce the probability of being fined if visited. y_t is a Bernoulli random variable indicating whether there was an inspection in period t, which occurs with a stationary probability $\theta > 0$. $Y_t = \sum_{s=1}^{t-1} y_s$ denotes the total number of visits until period $t-1$.

Updating of Beliefs θ is unknown to the vendor. She forms beliefs $\hat{\theta}_t$ about each day's visit probability on the basis of the history of visits $(y_1, \ldots, y_{t-1})$. We assume that the prior $\hat{\theta}_1$ is distributed Beta(α_0, α_1). Since y_t is Bernoulli, Bayesian updating implies $\hat{\theta}_t$ is distributed Beta$(\alpha_0 + Y_t, \alpha_1 + t - 1 - Y_t)$, and $\mathbb{E}[\hat{\theta}_t] = \frac{\alpha_0 + Y_t}{\alpha_0 + \alpha_1 + t - 1}$. The perceived probability increases with the *share* of periods in which the vendor has observed a visit in the past, adjusted by the strength of her prior (which is parameterized by α_0 and α_1).

Defensive Actions If a vendor decides to sell, she could either sell the hake *openly* or adopt costly defensive actions that reduce the probability of getting fined when inspected. If the vendor is inspected while selling *openly*, she is fined with probability one. The effectiveness of *defensive* actions depends on how knowledgeable vendors are about loopholes in the audit system. Vendors learn about enforcement loopholes as they receive visits.[11] We denote the probability of avoiding a fine through defensive actions $g : \mathbb{N}_0 \to (0, 1)$, where $g(Y_t)$ is a strictly increasing function of the past number of inspections.[12] We assume that the vendor can never make defensive actions completely foolproof, so $\lim_{Y \to \infty} g(Y) = \bar{g} < 1$.

Vendor's Problem In every period, the vendor chooses whether to sell hake openly, defensively, or not at all. $s_t = 1$ indicates the vendor sells hake in t, and $d_t = 1$ indicates the vendor adopts the costly defensive action. We solve the vendor's problem by backwards induction: Conditional on Y_t the vendor's expected utility from each type of selling strategy:

$$\begin{aligned} U[d_t = 0 | s_t = 1, Y_t] &= v - \Omega \mathbb{E}\left[\hat{\theta}_t\right]; \\ U[d_t = 1 | s_t = 1, Y_t] &= v - \Omega \mathbb{E}\left[\hat{\theta}_t\right](1 - g(Y_t)) - c. \end{aligned}$$

Proposition 1. *For any time t define the thresholds $\underline{\delta}_t = \frac{c}{\Omega g(Y_t)}$ and $\overline{\delta}_t = \frac{v - c}{\Omega(1 - g(Y_t))}$. Then*

- *When $g(Y_t) \le \frac{c}{v}$, the vendor never adopts defensive actions. She sells* openly *if $\mathbb{E}[\hat{\theta}_t] \le \frac{v}{\Omega}$, and does not sell if $\mathbb{E}[\hat{\theta}_t] > \frac{v}{\Omega}$.*

- *When $g(Y_t) > \frac{c}{v}$, the vendor sells hake* openly *if $\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t$; sells hake* defensively *if $\underline{\delta}_t < \mathbb{E}[\hat{\theta}_t] \le \overline{\delta}_t$; and does not sell hake if $\mathbb{E}[\hat{\theta}_t] > \overline{\delta}_t$.*

[11] For example, if the vendor observes that the inspector is reticent to levy a fine when she freezes the hake and claims it was caught in August, she will learn to adopt that strategy.

[12] This model assumes the learning takes place regardless of the action chosen by the vendor. A more sophisticated version could allow for action-dependent learning (bandit problem), which would add a dynamic component. Assuming the learning is independent of the action seems somewhat realistic in our context, and keeps the model simpler, preserving the key theoretical insights we can take to the data.

The proof of these results are in the Appendix. For Y_t high enough, $g(Y_t) > \frac{c}{v}$. As $g(\cdot)$ is increasing; (i) once $g(Y_t) > \frac{c}{v}$ this relation never reverses, and (ii) $\underline{\delta}_t$ is decreasing in Y_t, and $\overline{\delta}_t$ is increasing in Y_t. Together, this implies that if it becomes sensible for the vendor to adopt the costly defensive strategy in some period (given her beliefs), then that choice remains optimal for all subsequent periods.

The figure 3.1a provides a numerical example of the timing and scope of adoption of defensive actions. In the early periods, the vendor lacks experience to sell defensively. Once the vendor accumulates more experience (Y_t increases), adopting defensive actions becomes more likely, increasing the sale of hake. The figure 3.1b shows the overall sale of hake fish.

Short versus Long Run In the long-run (as $t \to \infty$), $\mathbb{E}\left[\hat{\theta}_t\right] \to \theta$, and vendor behavior is governed only by the structural parameters of the model, and the learning dynamics become irrelevant. Incentives to sell are lowered with higher visit intensity θ, higher long-run enforcement effectiveness in the presence of vendor adaptation $1 - \bar{g}$, and lower demand for hake v.[13]

Our modeling focuses mostly on the short run learning and adaptation, because these are the dynamics that we observe in our daily data collected during the hake ban in September. The short run comparative statics depend heavily on the specific form of learning and adaptation, $g(\cdot)$ and the vendor's prior belief (α_0, α_1). We focus on the most empirically relevant case for hake sales in Chile, in which the vendor's prior $\mathbb{E}[\theta_1] = \frac{\alpha_0}{\alpha_0+\alpha_1}$ is diffuse (i.e., $\alpha_1 \gg \alpha_0$), and she does not know the loopholes in the audit system before receiving any monitoring visits from this novel program we implement (i.e. $g(0)$ is small).

Section 3.3 describes comparative statics of varying enforcement design in this setting. In particular, sections 3.3 and 3.3 discuss the implications of varying the frequency of enforcement visits and the predictability of the enforcement schedule in our field experiment.

Enforcement and Learning

To perform the short-run comparative statics, we use the notation $\Delta x_t = x_t - x_{t-1}$ for any variable x, to define the effect of increasing visit frequency:

$$\begin{aligned}\Delta\left(\mathbb{E}[\hat{\theta}_t](1 - g(Y_t))\right) &= (1 - g(Y_{t-1})) \cdot \Delta\mathbb{E}[\hat{\theta}_t] - \mathbb{E}[\hat{\theta}_{t-1}] \cdot \Delta g(Y_t) - \Delta\mathbb{E}[\hat{\theta}_t] \cdot \Delta g(Y_t) \\ &\approx (1 - g(Y_{t-1})) \cdot \Delta\mathbb{E}[\hat{\theta}_t] - \mathbb{E}[\hat{\theta}_{t-1}] \cdot \Delta g(Y_t) \qquad (3.1)\end{aligned}$$

There is a threshold for the number of visits $\bar{Y} \in \mathbb{N}_0$[14] such that increasing inspections beyond $\bar{Y}$ has ambiguous effects on the vendor's propensity to sell. A new visit increases the vendor's perceptions about the probability of future visits ($\Delta\mathbb{E}[\hat{\theta}_t]$), but also allows her to acquire skills to circumvent the fine ($\Delta g(Y_t)$ is weakly positive).[15] At high values of Y_t,

[13] The specific long-run conditions are discussed in Appendix C.4

[14] Defined as $g(Y) \le \frac{c}{v}$ if and only if $Y \le \bar{Y}$. Such a $\bar{Y}$ exists and is unique if learning is effective enough: $\bar{g} > cv$, and due to the fact that $g(\cdot)$ is increasing.

[15] At $Y_t < \bar{Y}$, the vendor has not yet learned enough and the defensive strategy is still ineffective, so extra visits only disincentivizes hake sales through updates on $\mathbb{E}[\hat{\theta}_t]$.

a new visit could inadvertently increase the vendor's ability to sell hake illegally. Figure 3.1b simulates the effect on overall sales, under specific parameter values. The propensity to sell decreases immediately after the introduction of enforcement, but increases thereafter as vendors learn how to circumvent the enforcement. We will examine these patterns using our daily data.

Effects of Frequency of Enforcement Visits

Increasing θ has two effects in equation (3.1): (a) the threshold $\bar{Y}$ is reached faster, and (b) $\mathbb{E}[\hat{\theta}_t]$ increases faster as well. Greater visit frequency gives the vendor more opportunities to learn how to avoid paying the fine if inspected. The relative effectiveness of high versus low frequency enforcement in the short-run will depend on the specific period when the comparison is made. Figure 3.2 numerically simulates the model for high and low frequency of enforcement over 30 periods, under specific parametric assumptions described in Appendix C.5. Figures C.5.2a and C.5.2b plot the vendor's adoption of defensive strategies under those high and low frequency enforcement scenarios. More intense enforcement initially reduces hake sales faster (as vendors update more quickly about θ), but vendors also start adopting defensive actions earlier. This makes high frequency enforcement relatively less effective in later periods.

Predictability of Enforcement Visits

Vendors set up in different *ferias* on different days of the week, as described in section 3.2. If auditors focus enforcement efforts in a single feria within a circuit, or on the same day of the week, then their visit schedule becomes predictable. For simplicity, we assume that the circuit rotates between two ferias f^i, $i = 1, 2$, and in each period the vendor has the option to sell once in each of them.[16] At the beginning of each period, the vendor decides whether to sell in each of the ferias. Beliefs about the likelihood of a visit θ_t now needs a superscript θ_t^i ($i = 1, 2$), where i identifies each of two ferias. The vendor updates her beliefs about the probability of a visit in each feria by looking only at the history of visits at that feria. Appendix C.4 details why this corresponds to an optimal belief formation process. We define predictability of the auditing schedule as follows:

Definition 1. *A policy is* ***predictable*** *or* ***targeted*** *if either $\theta^1 = 0$ or $\theta^2 = 0$. A predictable policy* ***targets*** *feria i if $\theta^{-i} = 0$. A policy is* ***unpredictable*** *if $\theta^1 = \theta^2$.*

Proposition 2. *Define a fixed enforcement capacity $\Theta = \theta^1 + \theta^2$. When Θ is large enough, the most effective policy in the long run is the unique unpredictable policy $\theta^1 = \theta^2$, because that deters sales in both ferias.*[17]

[16] Modeling one feria per period (say, f^1 in odd and f^2 in even periods) yields the same qualitative insights.

[17] This is true for $\Theta \geq 2\overline{\delta}_\infty$, where $\overline{\delta}_\infty = \lim_{Y_t \to \infty} \overline{\delta}_t = \frac{v-c}{\Omega(1-\bar{g})}$ (see Proposition 1). At lower enforcement capacity, the regulator might do better by targeting a single feria, as explained in Appendix C.4

With fixed enforcement capacity, learning occurs at the same rate under either a targeted or an unpredictable policy. However, vendors are more likely to adopt defensive actions in the targeted feria under the predictable policy, while the probability of selling in the non-targeted feria inevitably will tend to one. We simulate the effects of predictable and unpredictable policies on hake sales in each feria in Figure 3.3a. Vendor selling strategies diverge between the targeted and non-targeted ferias, and under most functional forms, the unpredictable policy is more effective on average. Sales fall sharply immediately after the introduction of the enforcement, but this is only true for the targeted feria under the predictable policy. This is why the predictable policy is less effective overall.

Model Predictions

The model provides specific empirical predictions that we can use our field experiment to test:

(1) Increasing the frequency of enforcement has ambiguous effects on sales in the short-run.

(2) Predictable enforcement is less effective in the short-run.

(3) The probability of selling is not stable, but varies over time as vendors learn and adapt.

(4) Vendors shift sales away from ferias targeted by enforcement schedule to non-targeted.

(5) Vendors exposed to enforcement will learn and adopt defensive actions after a few periods.

(6) The information campaign reduces hake sells both in the short-run and the long-run.

3.4 Data

We deployed "mystery shoppers" to surreptitiously gather information about hake availability in fish markets, once during the ban (September 2015) before and after our interventions, and again six months later in March 2016. We conducted two rounds of surveys of consumers during those same two periods. We also surveyed fishermen at *caletas* and vendors at *ferias* to map the fish supply chain and investigate spillovers. Figure 3.4 describes the timing of the interventions and data collection activities. In total, seven different data sources are used in the analysis.

Mystery Shopper Surveys

For us to reliably measure illegal activity, fish vendors cannot know that they are monitored. This poses an interesting data collection challenge. To develop a strategy to address this challenge, the research team visited dozens of ferias before the ban to understand the market structure and relationships between vendors and consumers. We learned that vendors do not know most shoppers, so an unfamiliar face will not necessarily raise suspicion. This made it a good environment to deploy *mystery shoppers* and collect data surreptitiously. 29 enumerators were trained to work as mystery shoppers. They were mostly women between the ages of 40 and 50, because this demographic group represents the typical feria customer profile. The mystery shoppers were trained to look and act like ordinary shoppers, to pose as buyers and (try to) purchase hake fish from the vendors. Mystery shoppers were not told the treatment status of any market, to guard against the possibility that they inadvertently behaved differently in treatment and control markets. They visited each circuit three times on average during the September ban. We conducted an additional round of mystery shopper visits in March 2016 to track longer term effects outside the intervention period.

These mystery shoppers gathered information on whether it was possible to buy hake, and the market price of the fish. They also collected information on what else was available for sale at the fish stalls and their prices, and to note down what was being purchased by other shoppers in their presence. The visit protocol was piloted and refined through multiple pre-period visits to ferias to make sure we elicited the required information without raising suspicion. Given this methodology, we could not collect information about the total quantity of hake being sold, because that would be unnatural for a typical shopper to ask about, and it would have made the vendors suspicious. The main outcome variable that this survey therefore produces is an indicator for whether it was possible to buy hake at any particular stall. The mystery shoppers also noted down general characteristics of the stall and vendor. They also wrote down notes on the behavior of fish vendors, including conversations occurring in their presence. This is how we learned about the practice of selling "frozen hake", where the vendor kept the fish on ice and claimed that it was caught legally in August. Many of those same vendors admitted to our mystery shoppers that the "frozen" fish was in fact, fresh.

Identifying Defensive Strategies

Defensive strategies are at the heart of our theory on learning and adaptation. These are normally difficult to observe because they are illegal and designed to be hidden. However, our data collection strategy was designed to uncover such hidden actions. Mystery shoppers uncovered two strategies most commonly used by vendors to circumvent enforcement: They hide the hake they sell (instead of displaying it openly), and they put the fish over ice and claim that it was caught legally in August, and frozen since then. There are other possible illicit reactions that are impossible for mystery shoppers to observe safely, such as bribes paid or threats issued during vendor-inspector interactions.

Hiding: Mystery shoppers were trained to ask vendors for hake even if it was not visibly on sale in the stall. They noted down each occurrence of "hidden hake", but we never shared the specific vendor or feria identity with our government partners, so as to protect vendor privacy and abide by our research ethics protocol. These data were very useful for the evaluation, but were never used to target enforcement.

The hidden hake fish was often stored in a cooler behind the board that displayed the stall's fish prices. This is costly for vendors, because displaying the fish available for sale and attracting customers' attention are the main marketing tools at the vendors' disposal. Many of our mystery shoppers noted down in survey instruments that they observed regular consumers asking vendors for hake when it was not visible. The hiding strategy evidently works because many consumers are willing to partake.

Freezing: On paper, vendors are not allowed to sell hake fish in any form in September. In practice, Sernapesca inspectors were more lenient with vendors who were detected selling "frozen" hake. This is the practice of freezing the fish on ice and claiming that it was harvested in August, before the ban. We had not anticipated this reaction, but a couple of our mystery shoppers noted the practice for us early enough such that we were able to collect systematic data on it. Matching our mystery shopper data at the daily level to the administrative data on fines levied (from Sernapesca's registry of inspector visits) suggests that inspectors were much less likely to levy penalties when the vendor was claiming to sell "frozen" hake.

Selling frozen hake is costly for vendors because consumers prefer the taste of fresh fish, and because freezing requires freezers and access to electricity. Using our other rounds of data, we see that freezing is virtually non-existent during the rest of the year. So this does appear to be a strategy that vendors use to circumvent the September ban.

Consumer Surveys at Fish Markets

We also surveyed consumers before and after the ban period. A separate team of enumerators (distinct from our mystery shoppers) stopped consumers close the points of entry and exit for the fish market, and asked questions with a survey instrument in hand. To encourage unbiased responses, enumerators informed consumers that the survey was conducted by university-based researchers, and that it aimed to gather information about food consumption in ferias. They were not asked to provide any personal identifiable information, and we only inquired about the list of food purchased in the feria in the past month - avoiding asking direct questions about the consumption of hake. We also asked consumers to provide a sense of their home location on a physical map we carried, so that we could match their residence to the neighborhoods assigned to the information treatment.

In total, 3,300 consumers were surveyed in October 2015 through 54 enumerator visits, and 3600 in March 2016 through 95 enumerator visits. This produces two rounds of a repeated cross-section; the same consumers were not followed over time.

Survey of Vendors at Markets and Fishermen at Fishing Villages

We surveyed fish vendors in every market in our sample in June 2016, which is outside the hake ban period. We asked vendors about the suppliers and intermediaries they source their fish from, so that we could map out the supply chain. We also asked vendors about their contacts with fish vendors who operate in other circuits, in order to study spillover and network effects.

To understand whether the effects of our interventions were transmitted upstream via the supply chain, we conducted a survey of fishermen during July-August 2016 in every coastal village in the region where hake fish is caught and distributed. We surveyed 231 fishermen from 74 fishing villages (caletas). Figure C.1.4 in Appendix C.1 contains a map of all caletas and fish markets.

Surveying fishermen was valuable for two reasons. First, the interventions were designed to ultimately reduce illegal fishing, so understanding the activities of the fishermen is essential for public policy. Second, the treatment effects may have spilled over to control areas if treatment and control markets are served by the same fishing village. Understanding these supply-chain connections are important for analyzing spillovers. Figure C.1.8 organizes our interventions and data collection activities along the supply chain for fish.

3.5 Results

We report experimental treatment effects first, and then use the daily data to test the model's predictions on learning and adaptation. We registered this trial in the AEA registry before data collection was completed. Our approach to analysis and the outcome variables we focus on closely mirror the project narrative we uploaded before we had access to any data. We highlight the most notable departures from the pre-analysis plan (PAP) in Appendix section C.7. We did not delve into the details of testing a model of learning in the PAP. The various treatment arms appear well balanced in terms of baseline socio-economic and weather characteristics. Details are provided in Appendix C.3.

Empirical Strategy

Mystery shoppers visited several stalls in each market multiple times in September 2015. These visits created a stall-day level panel dataset of 906 visits. The first enforcement visit to various markets by Sernapesca officers occurred between Sept 4 and 10. Our panel data consists of 242 visits during the pre-enforcement period, plus 664 visits during the post-enforcement period. We use the following regression specification to evaluate the interventions, where each observation refers to a mystery shopper visit at fish stall s, in feria f, from circuit c visited on day t:

$$y_{sfct} = \beta_0 Post_t + \beta_1 T_c + \beta_2 T_c \times Post_t + \beta_3 y_{sfc0} + X_{ct}{}' \beta_4 + \varepsilon_{sfct} \tag{3.2}$$

y_{sfct} is the outcome variable, such as an indicator for whether illegal hake fish was available at that stall on that day. The treatment assignment (T_c) varies at the circuit level. The variable $Post_t$ indicates the post-intervention period, September 8-30.[18] We control for weather on each day, whether the inspector visited the market that day, a few socioeconomic covariates (e.g. municipality crime rate), randomization strata fixed effects, and the baseline (pre-intervention) value of the dependent variable. The error term, ε_{sfct}, is clustered at the circuit level, which was the unit of randomization. The coefficient of interest for the evaluation is the parameter β_2, which captures the difference between treatment and control groups during the post-intervention period. In most of our tables, we will only report the β_2 coefficients, and suppress all others.

To study consumer fish purchase behavior, we use surveys of consumers conducted at ferias. We use the following regression specification to evaluate the effect of interventions, where each observation refers to a to a single consumer i, surveyed in feria f, from circuit c:

$$y_{ifc} = \gamma_1 T_c + X_{ic}{}'\Gamma + \epsilon_{ifc} \tag{3.3}$$

Where y_{ifc} is the outcome variable, such as the number of times the consumer purchased hake fish in the past month. T_c is the treatment status at the circuit level, and X_{ic} represents a set of covariates, including socioeconomic characteristics of the municipality, individual demographics (usual fish consumption, age, gender, and household income) and strata fixed effects. Consumers are assigned treatment status based on the feria where they were interviewed.[19] Standard errors are clustered at the circuit level.

Treatment Effects on Hake Sales Observed by Mystery Shoppers

Column 1 of Table 3.2 shows the effect of the interventions on whether fresh, visible hake was available for sale in that stall, as detected by mystery shoppers. Column 2 shows effects on whether hake in any form (fresh and visible, hidden in the back, or "frozen" hake that is kept on ice) was available for sale. Each dependent variable is binary, and we report marginal effects from a Probit regression. The table 3.2 presents the coefficients of interest of regression equation 3.2, which track the effects of the demand-side information campaign, the supply-side enforcement treatment, or the interaction between the two (ferias where both interventions were simultaneously administered), during the post-intervention period.[20]

[18] Many of the information campaign letters arrived at households even after September 8. There are other reasonable ways to define the post-intervention period, and we make a conservative choice. We have verified that the exact definition of the post intervention period does not affect our main results.

[19] While that is the only sensible choice for the enforcement treatment, we could have also used the person's address to link them to the information treatment. Results look very similar either way, and we have imperfect information on individual addresses, so we use the feria location.

[20] We randomized the Information Campaign over the subset of the 48 most populous municipalities in our sample (out of 70 total). We control for an indicator for these 48 municipalities in all our regressions. We have also run regressions restricting the analysis sample to these 48 municipalities, and the results look very similar.

In column 1, vendors in markets exposed to the information campaign are 13.3 percentage points less likely to be selling fresh, visible hake relative to control group vendors.[21] This is quite a large effect, considering that about 43% of vendors in control markets were selling hake before the interventions were launched. Vendors operating in markets where Sernapesca monitors visit to levy penalties become 17.8 percentage points less likely to sell fresh, visible hake. The combination of the two treatments also produces a 17.9 percentage point decrease in hake availability, so there is no evidence that the information campaign complements the enforcement strategy to make it more effective.

When we add "hidden" and "frozen" hake to fresh/visible hake sales in column 2 to create a broader dependent variable that captures any type of hake sales, the enforcement treatment effects become smaller and lose statistical significance. Taken together, the two columns suggest that while the interventions reduced vendors' propensity to engage in illegal activity that could be easily monitored by regulators (visible sales in column 1), it is not so clear whether it actually reduced the underlying environmental harm that we care about (column 2). The reduction in the size of the treatment effect moving from column 1 to 2 stems from the defensive strategies that vendors adopt in response to the audits.[22]

Consumer Behavior

We consider the mystery shopper data to provide the most reliable measure of illegal behavior, but the consumer surveys at markets allows us to report changes in purchase patterns during the ban. The first column of Table 3.3 shows treatment effects on the number of times that consumers report buying hake fish during the previous month. The reported coefficients are marginal effects from a Poisson regression, evaluated at the mean of all covariates. We see significant decreases in (self-reported) hake purchase across all treatment arms, and so results are generally consistent with the mystery shopper survey. However, in these consumer reports, the treatment effects appear larger in information campaign areas where purchases decrease by 50% compared to the control group. This may be because the

[21] The "Information Campaign" group is a marker for circuits located in municipalities assigned to receive the High-Saturation Information Campaign, where the majority of neighborhoods were treated with the campaign. Appendix Table C.3.4 explains why we made this modeling choice. Our consumer survey data indicates that the majority (69%) of shoppers we found shopping at *ferias* located in "control" neighborhoods in high-saturation treatment municipalities resided in neighborhoods that were treated. It therefore makes more sense to code such *ferias* as 'treated' with the information campaign. Appendix C.3 shows the results of re-estimating the results in Tables 3.2, but reverting to coding ferias in control neighborhoods as not treated with information. The results are qualitatively similar. The high-saturation information treatment has significantly larger effects on hake sales than the low-saturation treatment.

[22] It is curious that the control group experienced larger reductions in "any hake" (column 2) than in "fresh, visible hake" (column 1). This is because a few control group vendors practiced freezing during the pre-intervention period (first week of September), but they stopped doing so after the interventions started. Apparently vendors in the control group learnt that there would *not* be much enforcement in their *ferias*, and reacted accordingly.

direct communication consumers received through the information treatment created some self-reporting bias.

Consumer behavior was also indirectly influenced by the enforcement activity. Not only did self-reported hake purchases decrease there relative to control markets, the third column also shows that consumers were about twice as likely (or 8-11 percentage points more likely) to mention to our enumerators, totally unprompted, that they did not buy hake because there was a September ban in place. Our enumerators did not ask consumers any questions about the ban, but were instructed to note down whenever a consumer spontaneously mentioned the ban. Consumers treated with the information campaign were 15 percentage points more likely to mention the September ban unprompted, so evidently the treatments were at least successful in spreading more information and awareness relative to control areas.

Variations in the Design of the Enforcement Strategy

We experimentally manipulated the enforcement schedule in two dimensions: Predictability and Frequency. Table 3.4 uses the mystery shopper data, and repeats the regression setup of Table 3.2, except that the enforcement treatment is now sub-divided into areas where the monitoring schedule was either predictable or unpredictable (column 1), or sub-divided into areas where monitoring was conducted at high versus low frequency (column 2). These provide direct tests of Predictions **(1)** and **(2)** highlighted in Section 3.3.

The first column shows that the enforcement strategy was more effective when it was unpredictable. When enforcement follows a predictable schedule (e.g. every Tuesday at 10am), its effect is not statistically different than zero. However, when we make the monitoring visit schedule difficult for vendors to predict, we see that there is a much larger and statistically significant decrease of 19 percentage points in vendors' propensity to sell hake, even after we account for vendor defensive reactions like hiding and freezing. The effect of the unpredictable schedule is statistically significantly larger than predictable enforcement. The lack of predictability makes it difficult for vendors to anticipate the visit pattern and deploy effective defense.

The second column shows results separately for the subgroup of vendors who received monitoring visits once a week (low frequency), and other vendors who were visited twice a week, which means that monitors followed a circuit around in the different market locations where those vendors set up stalls on different days of the week (high frequency). The high frequency visits in principle limit opportunities for spatial and temporal displacement of illegal hake sales. The strategy of devoting additional resources to enforce at high frequency backfired. Enforcement is more effective at reducing hake availability in markets that were visited less frequently. Although the 9.2 percentage point gap between low and high frequency is meaningful in magnitude, it is not statistically significant.

Evidence on the Process of Learning and Adaptation

Number of Visits: In this section, we study some of the specific theoretical predictions on how vendors learn and adapt to enforcement by merging our daily data collected via mystery shoppers with the administrative data from Sernapesca inspectors. Observations made by mystery shoppers at a specific feria on a given day are linked to the history of enforcement visits in that feria and circuit. Appendix C.3 describes Sernapesca's enforcement activities in more detail. We organize the data this way to test Prediction (3) from Section 3.3 that vendors' propensity to sell would not remain stable as they learn about enforcement and adapt.

Figure 3.5 compares the week-to-week behavior of vendors exposed to different frequencies of enforcement. Consistent with the theory of learning, the two treatments produce similar effects at the beginning of the month, but effects diverge over time. Vendors exposed to higher visit frequency sell *more* hake, not less by the end of the month. This is consistent with the idea that more interactions with auditors allow vendors to learn about enforcement loopholes and adopt defensive strategies.

Figure 3.6 plots the likelihood of selling hake on a given day as a function of the number of inspections received at that *feria* until that day.[23] We see that receiving more visits reduces the probability of selling over time. However, the effect is non-linear: Earlier visits have a larger effect on reductions in hake sales than subsequent visits. This is especially true in the experimental arm with a predictable visit schedule. This differential effectiveness over time in the predictable arm was also evident in the theoretical simulations [Figures C.5.4a and C.5.4b]. The intuition is that vendors learn that one of the ferias where they sell is not being targeted, and continue selling illegally at that location.

Schedule of Visits: We test this intuition directly in Appendix C.3 using *within-circuit* variation to study whether the same vendor shifts sales to non-targeted days and markets. Each fish vendor rotates between ferias within a circuit on different days of the week in a pre-determined pattern. We see substantial evidence supporting the this form of day-of-week displacement (model Prediction **(4)**).

In Appendix Table C.3.10, vendors who experienced inspections in different ferias on different days of the week (DOWs) reduce hake sales by an extra 9 percentage points (p-val¡0.01) in the second half of the month, relative to vendors who were targeted at a single feria always on the same DOW. Furthermore, Tables C.3.11 and C.3.12 study vendors' decisions to sell hake in the *non-targeted* feria in the second half of the month in circuit-fixed-effects regressions. We see that the same vendor sells more at markets and weekdays

[23] The estimates are obtained from the following regression specification:

$$Y_{sfct} = \sum_{n=0}^{N} \left(\beta_n^P \times \mathbf{1}(\#\text{Enf}_{ct} = n) \times Pred_c + \beta_n^U \times \mathbf{1}(\#\text{Enf}_{ct} = n) \times UnPred_c \right) + X_{ct}{}'\Gamma + \varepsilon_{sfct} \qquad (3.4)$$

The term $\mathbf{1}(\#\text{Enf}_{ct} = i)$ indicates circuits that have been visited n times by Sernapesca officials at the moment the secret shopper collected the data.

where she did *not* experience a visit, relative to another market/weekday where she did, and is also more likely to shut down her stall entirely in the "targeted" market.

This behavior closely resembles the theoretical simulations displayed in Figure 3.3, where we explain how the shift in sales towards non-targeted ferias results in the relative ineffectiveness of the predictable schedule of enforcement. We relegate this "within-circuit" evidence to the appendix, because Sernapesca chose which feria to visit within each circuit partly based on logistical considerations, and this variation therefore cannot be treated as random.

Defensive Strategies We now study defensive strategies highlighted in Prediction **(5)** in section 3.3. Our mystery shoppers collected systematic data on vendors' propensity to sell "hidden" fish from the back that was not displayed at the stall, and "frozen" fish that they claimed was caught in August. Hiding was clearly used as a defensive strategy to circumvent the September ban: we conducted another mystery shopper survey six months after the ban, and we did not observe even a single stall selling fish that was not publicly visible at that time. There are several pieces of circumstantial evidence in our data that freezing is also pretense; that fishermen and vendors are not actually protecting the environment by catching fish in August and freezing it until September. First, we document more freezing in the second half of September 2015 than during the first half, after vendors have had a chance to learn about the enhanced regulatory activities. Real freezing would have been much less costly to engage in during the first half of the month. Second, we collected data on stall characteristics, and availability of a freezer in a stall is not at all predictive of freezing. If anything, our mystery shoppers find that stalls without freezers are more likely to be selling frozen fish post-intervention. Third, many secret shoppers noted down that in their conversations with vendors, many vendors admitted (and even insisted) that the fish was fresh even though it was labeled as frozen.

Figure 3.7 shows the prevalence of freezing and hiding across treatment groups. We divide up the control group into markets that have another circuit that is randomly assigned to enforcement within 10 kilometers (to capture any information spillovers), and *pure control* markets that are more than 10km away from any treated area. Several notable patterns emerge:

1. We do not observe any hiding or freezing at all in pure control markets in the post intervention period. In contrast, 7.2% vendors operating in circuits that received Sernapesca inspector visits sell frozen fish (p-value ¡0.01), and 3.2% of those vendors engage in hidden hake sales (p-value 0.01).

2. Vendors operating in circuits exposed only to the information campaign did not engage in any hiding or freezing at all. Vendors (sensibly) employ these defensive strategies only against Sernapesca inspectors, not informed consumers. Evidently there is something fundamentally different about targeting the demand side: The information campaign did not simply signal enhanced government attention to the problem. The

consumers are an important independent actor whose knowledge affects vendor behavior. This result also explains why the enforcement strategy appeared to produce larger decreases in fresh, visible hake sales than the information campaign (column 1 of table 3.2), but not once you take vendor defensive strategies into account (column 2).

3. 4% of vendors who operate in control markets - but located close to treated areas - engage in hiding and freezing, in contrast to 0% in *pure control* markets (p-value 0.02). There appear to be some spatial spillovers in information about Sernapesca visits, and in vendor behavior. We will explore these spillovers at greater depth in Section 3.6.

Figure C.2 shows that the proportion of hake vendors who adopt these defensive strategies increase week-to-week in response to the enforcement activities. Figure 3.8 describes the proportion of all hake sold in frozen or hidden form (as opposed to fresh, visible form), within the set of stalls that sell hake at all. Freezing and hiding were extremely unusual at the beginning of the month. Vendors exposed to Sernapesca enforcement increasingly adopted these defensive strategies week to week, vendors not exposed to enforcement did not. By the end of the month, nearly 70% of the stalls selling hake in the enforcement areas hid or froze. Consistent with model Prediction **(5)**, these data strongly suggest that these were indeed defensive strategies employed in response to enforcement.

Finally, consistent with Prediction (6), the information campaign reduces hake sales even after accounting for defensive strategies (as already shown in Table 3.2), and consumer surveys conducted again 6 months after the ban ends (see table C.3.14) shows that the demand-side effects somewhat persist over the long run.

3.6 Spillovers and Market Level Effects

While our experiment was targeted to reduce hake sales in treated *ferias*, it may have had spillover effects on control markets through information transmission, or by changing equilibrium prices (Blattman et al., 2017). It may also have affected the behaviors of other market actors, such as the fishermen who supply to vendors. It could have also changed the prices and quantities of other fish that can act as substitutes for hake. We collected additional data to study these spillovers and equilibrium effects, including a survey of fishermen, a survey of vendors to understand their social and supply-chain connections to vendors operating in other markets, GIS data on the location of all markets, and data on the prices and availability of substitute fish. The vendor and fishermen surveys allow us to map the supply chain for each of the ferias in our sample. The geography of Chile (with a very long coast) creates large spatial variation in the locations of *ferias* where vendors sell and *caletas* where the fishermen bring in their catch, which in turn produces variation in geographic and social connections between different market actors (see Figure C.1.4).

Spillovers on Control Markets

We identified three primary channels through which treatment may affect behavior of control markets, and collected data on each channel:

1. *Spatial spillover:* Control markets located geographically close to a treated market may feel the effects of treatment because they share consumers with the treated area.

2. *Social spillover:* If control market vendors are socially connected to vendors operating in treatment areas, they may be more likely to learn about *Sernapesca*'s enforcement activities.

3. *Supply chain spillover:* Treatment and control vendors may source from the same fishermen. If a supplier changes fishing behavior due to treatment, that could indirectly affect fish sales in control markets.

Of these different channels, an increase in fish sales in control markets due to indirect effects is of greatest econometric concern. If fishermen dump all excess hake in control markets when vendors in treated markets are unwilling to buy hake, then the treatment-control difference will appear to show that the treatment was effective, when in fact hake sales were simply spatially displaced towards the control group. Our regressions would over-estimate the effects of treatment in that scenario. This is why it's important for us to re-investigate these effects controlling for these sources of spillovers.

In Table 3.5, we re-estimate the effects of predictable and un-predictable enforcement originally reported in Table 3.4, but now controlling for potential channels of spillover effects.[24] The main treatment effects get a little stronger after controlling for spillovers, but the spillover effects are only suggestive and statistically imprecise.

The first column presents the benchmark: unpredictable enforcement reduces hake availability by 15.7 percentage points in this specification without accounting for any spillover. The second column controls for spatial spillovers, with the indicator "within 10km of Treated Market" turning on for untreated markets that have at least one treated *feria* within a

[24] We follow a procedure similar to Miguel and Kremer (2004) in estimating treatment effects in the presence of spillovers. We divide the control markets into subgroups; (a) Control areas that are more likely to have been affected by treatment due to geographic or social or supply chain connections, which we call "Spillover Group", and (b) Control areas un-connected to treatment markets, which we call "Pure Control". Note that sub-dividing the control group this way reduces the number of markets allocated to the omitted category. To retain sufficient statistical power, we therefore focus on re-estimating the effects of enforcement treatment variations only, because spillovers cause the greatest econometric concern (of over-estimating treatment effects) for this particular result. In this setup, some of the markets in the omitted category received the information treatment, so the regression coefficients will look a little smaller in this table compared to Table 3.4. For the same statistical power reasons, we only study an overall spillover effect of enforcement, and do not try to estimate separate sub-treatment spillovers.

distance of 10 kilometers.[25] [26] The coefficient of this variable is negative but small and statistically indistinguishable from zero, suggesting very limited spillovers based on shared consumers due to geographic proximity. The third column includes an indicator for control markets where at least one vendor reported that they knew a vendor in a different market that was randomly assigned to the enforcement treatment. The coefficient on this variable suggests that there was a 7 percentage reduction in hake availability in markets experiencing this "social spillover", but this effect cannot be statistically distinguished from a null effect with any confidence. Controlling for this form of spillover increases the effect of unpredictable enforcement to a 19.9 percentage point reduction in hake availability ($p < 0.05$). Finally, column 4 includes an indicator for control markets who source from fishermen operating out of *caletas* that primarily supply to other markets that were assigned to the enforcement treatment. We see a 7.7 percentage point reduction in vendors' propensity to sell hake in control markets that are connected to treated markets through shared suppliers, but the effect is again not statistically precise.

Importantly, accounting for these spillover effects make the main treatment effects of unpredictable enforcement on enforced areas a little larger and more statistically precise. This is because controlling for spillovers allow us to compare treated areas to the subset of "pure" control areas unaffected by the treatment.

Change in Number of Stalls Selling Fish

Our intervention may force some fish vendors to exit the market altogether. Appendix Figure C.2.2 shows that the average number of fish stalls decreases in the markets randomly assigned to the enforcement treatment, especially during the second half of September. This itself is an important effect of the treatment, but it is not captured by the treatment effects reported in Table 3.2. Table C.3.13 in Appendix C.3 describes how we correct our estimates for stalls exiting. The correction makes the effect of enforcement larger, but it does not affect the coefficients for other treatments very much.

Treatment Effect Transmission along the Supply Chain

For the supply chain spillover channel to be relevant, the fishermen supplying hake to these vendors must have altered their behavior in some way. To understand those changes, we directly survey fishermen operating out of every *caleta* (fishing village) that serves the mar-

[25] Using the 10 km radius evenly divides the control group into "pure control" and "spill-over market", and therefore maximizes statistical power. Alternative definitions produce similar results.

[26] Vendors connected to a larger number of other circuits are more prone to being exposed to the treatment, and that variation is not random. To control for this, we include a full set of dummy variables for the number of other circuits that each reference circuit is connected to, separately for spatial, social and supply-chain connections. Thus, the variation of exposure to spillovers stems only from the treatment status of other markets, which is exogenous because it was randomly assigned. (Miguel and Kremer, 2004).

kets in our sample.[27] The reactions of fishermen are particularly important to track because our interventions conducted at the final point-of-sale has to somehow get transmitted up the supply chain to fishermen, for these interventions to ultimately protect the hake population. Only if fishermen start perceiving the effects of these interventions on demand conditions will they change fishing behavior in ways that improve environmental outcomes.

Since we did not have baseline data from fishermen for years preceding the September 2015 ban, we ask them retrospective questions in 2016, in which the fishermen are asked to compare demand and profits during September 2015 (when our interventions were launched) relative to September 2014. To minimize possible response bias given the government fishing ban, we were careful to phrase our questions generically, to cover revenues earned from all types of fish, and not just hake specifically. Retrospective answers may be subject to recall bias, but since these fishermen were not directly treated, it is less likely that the recall bias is correlated with treatment assignment. To report treatment effects on fishermen, we have to connect each *caleta* to treatment and control markets. We use the vendor survey on the structure of the supply chain -i.e. which caletas each vendor buys from - to link fishermen to the randomized treatments.

Table 3.6 reports results. Column 1 shows that fishermen operating out of *caletas* that sell to at least one circuit which had been randomly assigned to enforcement, are 24 percentage points more likely to report that they earned less in September 2015 from all fishing activities compared to September 2014, relative to fishermen in caletas that supply to control group *ferias*.[28] Fishermen operating out of caletas that supply to both enforced markets and to markets that experienced the information campaign were 36 percentage points more likely to report lower revenues during the month of the interventions, compared to the same month in the previous year. Treatment effects were perceived by fishermen upstream in the fish supply chain. Column 2 shows that these fishermen are more likely to report that vendors were less willing to buy hake in September 2015 compared to the previous year, but this result is marginally significant with ($p < 0.10$). Column 3 shows suggestive evidence ($p < 0.10$) that fishermen linked to the information campaign areas are more likely to report that final consumers are aware of the hake ban.

Effects on Fish Substitutes

We collected data on prices and availability of other fish species in the same markets where hake is sold. The September ban is only specific to hake fish, so we might expect consumers to substitute to other fish varieties. This may be because informed consumers choose to avoid hake fish during the ban, or because the enforcement treatment reduces hake availability or increases its price.

[27] A few caletas in the regions covered by our sampling frame are only used by divers who harvest seafood, not fish -and we therefore exclude those *caletas*.

[28] We could instead define exposure based on the proportion of circuits enforced, and results look similar. The "at least one" formulation is attractive because this indicator evenly divides the sample into equal halves.

The universe of data from all markets suggests that there are seven possible fish substitutes for hake,[29] but a typical stall only offers two or three varieties of fish. Table C.3.2 in the Appendix describes the availability and price of different fish species observed by mystery shoppers in ferias during September 2015. The most common fish substitute is pomfret, which can be found in two-thirds of all markets. Pomfret is larger and (arguably) more tasty than hake fish and is not over-exploited. In Table 3.7, we study the availability of pompfret (column 1), or any other non-hake fish including pomfret (column 2), as a function of the treatment status of the market where the fish stall is located.

The penultimate row of the table indicates that stalls in control markets are 29 percentage points more likely to start selling pomfret during the September hake ban, so it appears that vendors in general move towards substitutes during the ban. The increase in pomfret sales during September is larger in treated areas (by a further 12-15 percentage points, which results in a 41-44 percentage point increase during the hake ban), but the treatment-control differences are barely statistically significant.[30] The p-value for only one of the three coefficients (associated with Predictable Enforcement) is below 0.10. Column 2 investigates treatment effects on the vendor's decision to offer each of seven different fish substitutes for hake. The sample size is larger in this regression because selling each fish variety is treated as a separate decision, but our standard errors are still clustered by the unit of randomization of the treatment (the circuit). The coefficients indicate that vendors who faced unpredictable enforcement become 15.5 percentage more likely to switch to selling other fish during the hake ban, compared to the 9 percentage point increase in control markets. This 6.5 percentage point treatment-control difference is statistically significant (p=0.051).

Effects on Prices

We collected data on fish prices during all our mystery shopper visits. However, prices are observed only when the fish is available for sale and hake is only available in 26% of markets. Treatment changes the propensity to sell illegal hake fish, so it affects the selection of which prices are observed. There are therefore large sample-selection issues that complicates any analysis of treatment effects on prices, and we refrain from running regressions on the price of hake. The most consumed fish during September (and second most consumed fish during the rest of the year) is Pomfret, which is available in 68% of the stalls (see Appendix Table C.3.2). Since pomfret is more often available (and not banned), we instead run regressions to study treatment effects on the price of pomfret.

[29] They are pomfret, mackerel, silverside, salmon, sawfish, albacore and southern hake. Of these substitutes, the southern hake is the only one with a similar ban, but in August. The southern hake is considerably larger than the common hake and is harvested in the southern regions of the country, without any geographical overlap with the common hake. More details are available in Subpesca (2015).

[30] Consumers are more prone to substitute products at similar price levels (see Table C.3.2). The hake is considerably cheaper than the pomfret and other relevant fish species. This fact may have limited the willingness to substitute for different fish species.

As a descriptive exercise, Figure C.2.3 shows that the price of hake increased week-to-week in September, over the course of the ban period. Pomfret prices fell by 10% in the second week and that lower price remained stable thereafter. This time-series pattern in prices is consistent with fishermen upstream in the supply chain shifting away from hake and towards catching pomfret during our interventions in September 2015. Through conversations with fishermen during our survey, we learned that they are able to adjust their fishing strategy to target different species if there are market signals that hake demand is low. To do so, they change the location and depth at which their nets are dropped.

Table 3.8 shows treatment effects on pomfret prices (column 1) and prices of all substitute fish including pomfret (column 2). We find that the price of substitutes weakly increase ($p-value < 0.1$) in markets where the information campaign discouraging hake consumption in surrounding neighborhoods, suggesting that part of the demand for hake shifted towards substitutes. Relative to the control group, markets that received enforcement show small and insignificant price decrease. The fact that we observe these differential price effects suggests that fish markets are at least somewhat segmented.

3.7 Cost-Effectiveness of Enforcement vs. Information

Given the complications associated with enforcing regulations documented in this paper, and the complexity of designing regulations that are robust to unanticipated defensive reactions from enforced agents, it is useful to determine how cost-effective the enforcement strategies were relative to an information campaign. We collected data from *Sernapesca* on the full administrative costs of implementing each treatment, so that we can report on the relative cost-effectiveness of enforcement and information strategies.

We define effectiveness of our interventions on the basis of our treatment effects on all hake sales (visible, hidden or frozen). Since the fish sold in ferias comes directly from fishermen villages and was harvested the same day or the day before, we assume that reduced hake sales is proportional to the decrease in hake fishing. The fishermen survey results reported in section 3.6 suggests that fishermen did feel the effects of the interventions. Our interventions were conducted at scale covering all major markets where hake is sold, which implies that our data are net of "leakages" of hake from our sampling areas.

In Table 3.9, we conduct the relative cost-effectiveness analysis by taking our best estimates of the effects of treatments on reduction in hake sales and combining it with an estimate of the number of fish available in the market that we compute using the data we collected from vendors. This allows us to create an estimate of the extra hake fish that are "saved" due to these treatments. Methodological details underlying these calculations are in Appendix C.6.

We compare this number with the cost of implementing each intervention to compute how much it cost to save each fish under each of the treatment assignments. Overall, the information campaign appears more cost effective than the enforcement strategy. This is partly because enforcement becomes less effective as vendors learn to hide and freeze fish and circumvent regulation. Enforcement costs US$6.05 per saved fish, compared to $4.98 under the information campaign.

However, once we examine specific versions of the enforcement strategy that were more successful at curbing hake sales, we see that sending monitors on an unpredictable schedule is a more cost effective way to protect hake, even after accounting for the fact that unpredictable monitoring schedules were more costly for *Sernapesca* to maintain because it required slack personnel capacity. The cost of "saving" a hake via unpredictable enforcement drops to $4.51. Not surprisingly, less frequent monitoring schedule is most cost-effective ($4.13 per saved hake) because it was both more effective at reducing hake sale than high-frequency enforcement, and it was obviously also cheaper to implement. Predictable and high-frequency audits were total policy failures in that they were 250-400% too expensive per hake saved, given the subversive adaptation by hake vendors.

These calculations are useful to gauge the *relative* cost-effectiveness of alternative strategies to protect hake, but it does not tell us whether any of these strategies would pass a cost-benefit test. Sophisticated benefit calculation would require us to take a stance on the biology of hake fish (how saving a hake in September 2015 translates into a dynamic effect on the hake population via reproduction), and the ecological value of protecting hake. These considerations are outside the scope of our analysis, but our results can be easily combined with benefit numbers from ecology studies. The analysis in this paper takes the government's regulatory goal ("Protect hake fish through a September fishing ban") as given, and studies the consequences of enforcing that regulation, and analyzes the best ways to achieve that goal.

3.8 Conclusion

Research in many fields of applied microeconomics evaluate the effects of new regulations, such as anti-corruption campaigns, fines for non-compliance with health, hygiene or environmental standards, or penalties for tax evaders. The effectiveness of such policies depend on the (sometimes unanticipated) reactions of the regulated agents to the new enforcement regime, which is in essence a micro version of the "Lucas critique" (Lucas, 1976). Agents adapt once they have had a chance to learn about the new rules, and may discover new methods to circumvent the rules. This paper presents a research strategy - composed of an experimental design and creative data collection - that permits an investigation of the effects of regulation net of agent adaptive behaviors.[31] This research approach should be broadly useful for policy evaluation whenever agents can adapt to circumvent enforcement. As one

[31] An alternative evaluation strategy would be to collect data in the short run before agents have an opportunity to react to the new regime, and in the long-run after they have reacted. This is more expensive, requires more time, and fundamentally more difficult, because researchers do not always know when and how agents would learn and adapt.

important example, such concerns were first-order in the design of the Dodd-Frank Wall Street Reform and Consumer Protection Act in 2010 following the global financial crisis. Smith and Muñiz-Fraticelli (2013) write about this regulatory effort:

> "[A] major problem with the new financial legislation is that it is responsive to past market innovations without being sensitive to future innovations (...) The problem is that these actors will not always behave in a predictable way. That is the genius of financial innovation; the market always looks for new opportunities for profit, and, as the dawn follows the dark, mischief may arise."

Our experimental variations that change the specific attributes of enforcement policy yield novel empirical insights about the adaptive behavior of regulated agents, and how to better design policy accounting for their adaptation. Data collected via mystery shoppers help us identify the ways in which agents exploit loopholes to continue selling fish illegally. Implementing a high frequency monitoring schedule produces a counter-intuitive result – but one that economic theory can rationalize – it allows vendors to learn the regulators' strategies faster, and more effectively cheat, thereby undermining enforcement efforts. Our theoretical and empirical results imply that even if monitoring is cheap, the regulator may do better by holding back some enforcement resources. And when you do monitor, adopting an unpredictable schedule makes it more difficult for agents to circumvent enforcement and proves to be the most cost-effective way to reduce hake sales even though it is more expensive to implement.

We use multiple surveys of different market actors to document that these interventions travel downstream to affect consumer behavior and travel upstream to affect the behavior of fishermen who supply to vendors. Our investigation of vendor reactions through mystery shoppers, spillover effects on other market actors, and benchmarking these results against the effects of an information campaign, all combine to produce a comprehensive evaluation of an important environmental program.

Ultimately we learn that without sophisticated design-thinking, attempts at enforcement can backfire. Designing and implementing a consumer information campaign is a much less complex task, it leverages consumer ethics (Hainmueller et al., 2015), and many regulators may rationally choose to proceed with such simpler approaches. After observing the results of this evaluation, the Chilean government decided to scale-up the information campaign during the 2016 ban on hake fish sales, and conduct similar information campaigns for fishing bans for three other species.[32] While the unpredictable, low-frequency monitoring proved to be the single-most cost-effective strategy in our evaluation, the government correctly surmised that vendors may have other second and third order subversive adaptations to audits in the long run. In contrast to an enforcement strategy which may need to be constantly revised in response to regulated agents' adaptation, the information campaign is easier to replicate

[32] See https://www.povertyactionlab.org/sites/default/files/documents/creating-a-culture-of-evidence-use-lessons-from-jpal-govt-partnerships-in-latin-america_english.pdf

and scale, especially once the government has already incurred the fixed costs of developing campaign materials.

Figures

Figure 3.1: Probability of Selling Hake

(a) Adoption of Defensive Actions

(b) Selling Decision

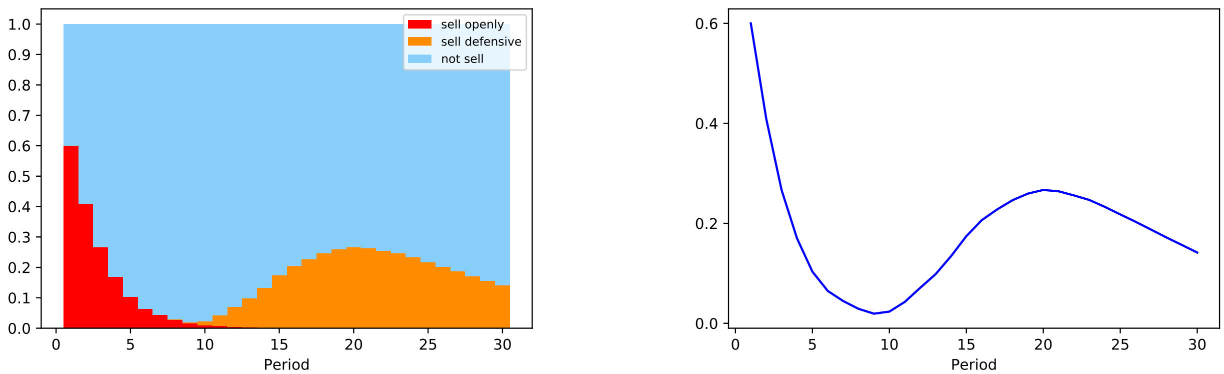

Notes: Figure C.5.2a and C.5.2b describe vendors' decision on whether and how to sell. This simulation uses the same parameters than previous graph: $\theta = 0.5, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + e^{-2\times Y + 12}\right)$, i.e.,$\bar{g} = 0.7$. The adoption of defensive strategies starts after a number of periods.

Figure 3.2: Probability of Selling Hake

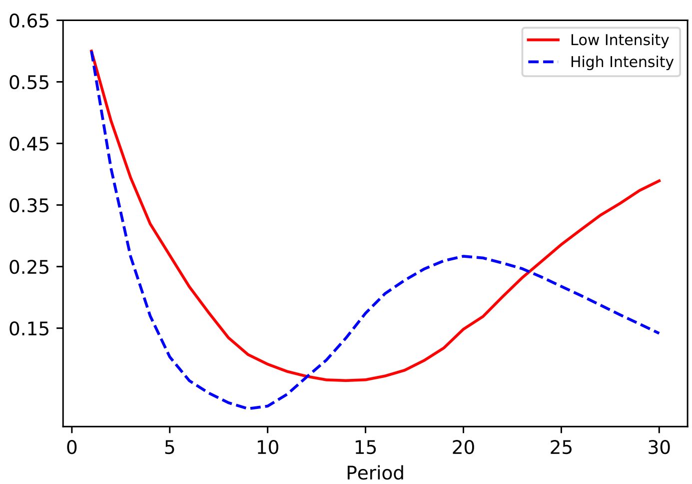

Notes: This figure shows the proportion of times in which a vendor sells hake depending on the frequency of the visits. This graph depicts 1000 simulations using the following model parameters $\theta^{high} = 0.5, \theta^{low} = 0.3, v \sim U(1/2, 3/2), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + exp\left\{-2 \cdot Y + 12\right\}\right)$, i.e., $\bar{g} = 0.7$. The probability of selling decreases quickly as the enforcement begins, however it increases as vendors learn about enforcement weaknesses. After a number of periods, it converges to the "long-run" equilibrium based on model's structural parameters.

Figure 3.3: Probability of Selling Hake

(a) Targeted vs. Non Targeted Ferias   (b) Predictable vs. Unpredictable Enforcement

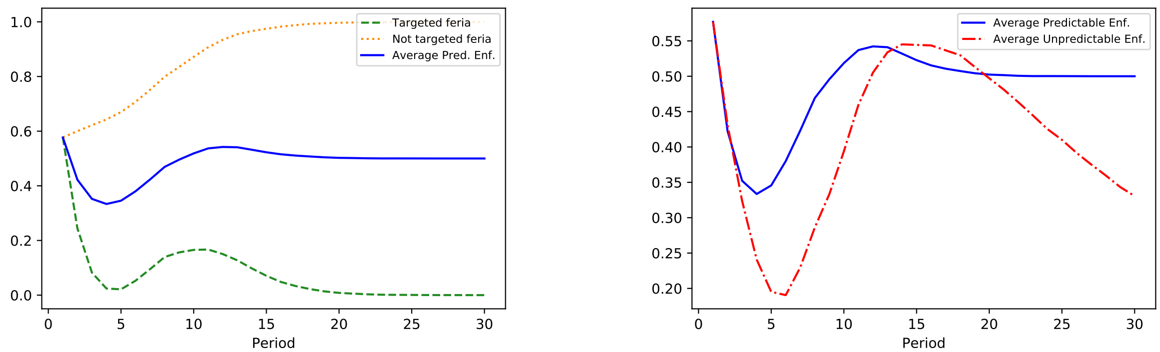

Notes: Figure 3.3a compares vendors' decision in targeted and non targeted ferias, assuming that vendors alternate between these ferias. The dashed line correspond to the average probability of sale, which is calculated assuming that in every period there's one half of the vendors in each type of feria. Figure 3.3b compares the average probability of selling under predictable vs. unpredictable enforcement. These simulations use the following model parameters: $\theta = 0.4, v \sim U(1/2, 3/2), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/(1 + exp\{-8 \cdot Y + 28\})$, i.e.,$\bar{g} = 0.7$.

Figure 3.4: Timeline of Interventions and Data Collection

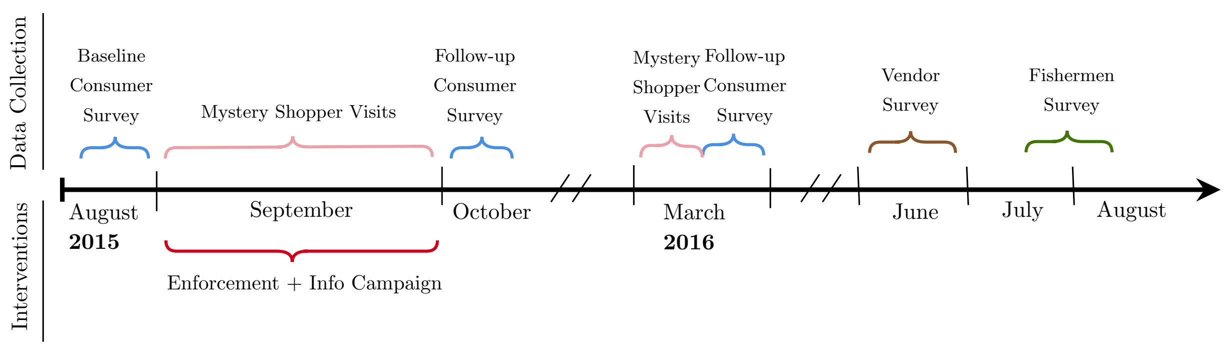

Figure 3.5: Hake Available

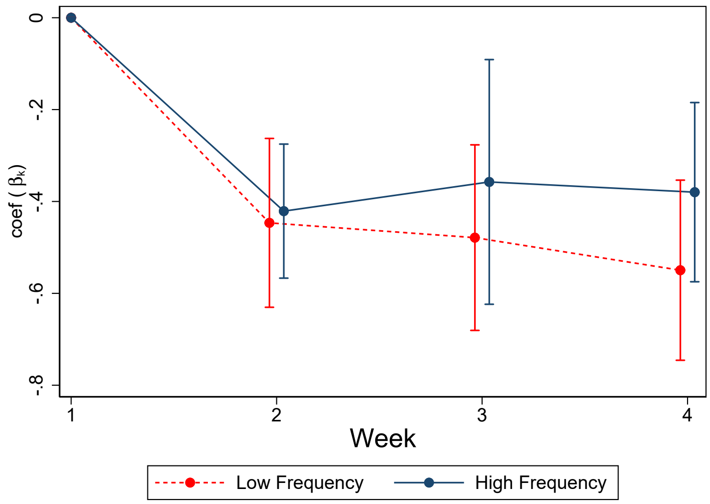

Notes: This figure shows how the sale of hake evolved week by week. The graph plots the coefficients of the treatment-week interactions. Each relevant coefficient is normalized relative to the first week. We exclude the first three days of the month to keep the weeks balanced, i.e., the first week starts on Sept 4th and ends on Sept 10th. Each regression controls for crime rate and strata fixed effects and the average outcome variable before the implementation. We cluster standard errors at the circuit level.

Figure 3.6: Hake Available

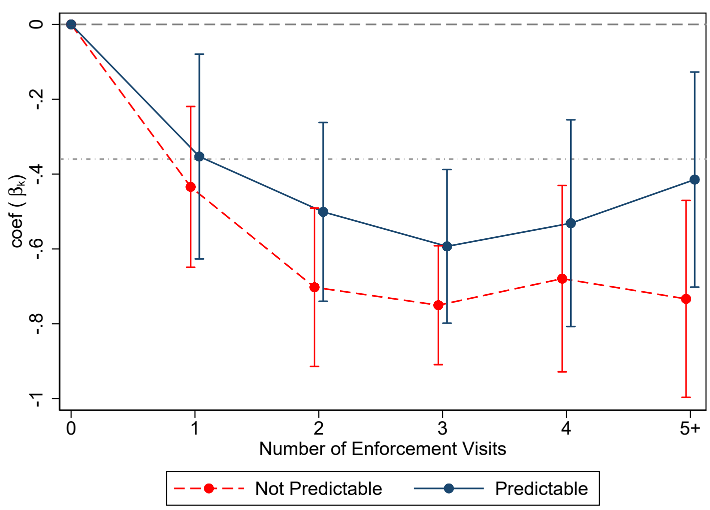

Notes: This figure shows how the sale of hake depends on the number of visits received until (including) the day the mystery shopper observed the behavior of the vendor. The horizontal line at -0.36 serves as a reference for the decrease in the probability of selling hake in the control group. This specification controls for crime rate, strata fixed effects, and the average outcome variable before the implementation. We cluster standard errors at the circuit level.

Figure 3.7: Hidden and Frozen Hake Fish

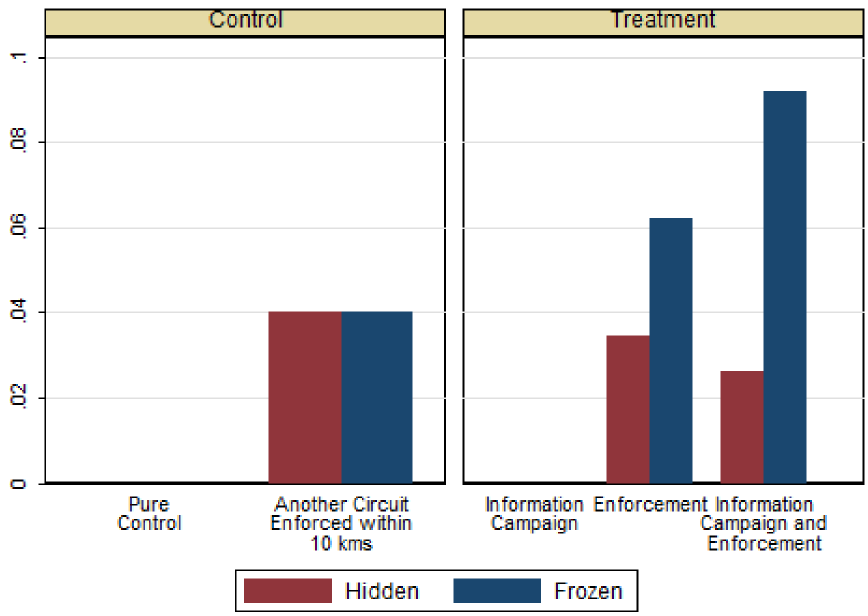

Notes: This figure shows the unconditional mean of hidden hake for different treatment status. The level of frozen hake is statistically different from zero for markets assigned to Enforcement and Enforcement and Info Campaign. The level of hidden is statistically different from zero for markets with Enforcement and spill-overs. Standard errors are not shown in the figure, but the accompanying text describes p-values of relevant comparisons.

Figure 3.8: Proportion of Defensive Hake

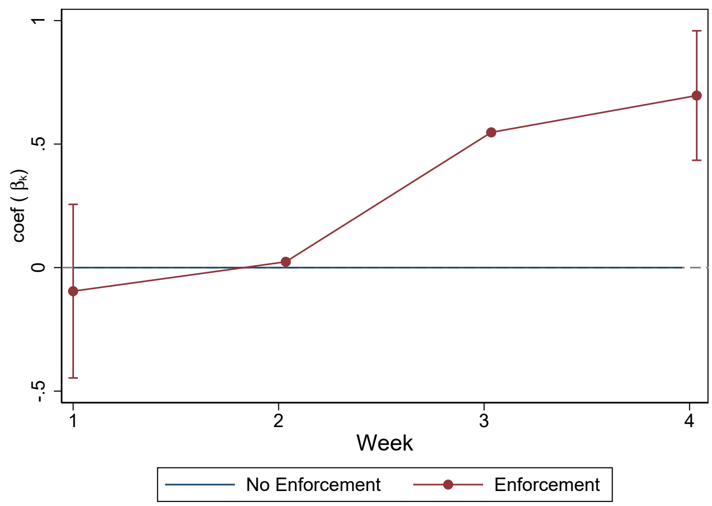

Notes: Figure 3.8 describes the conditional probability of selling hake either frozen or hidden. The coefficients were obtained from an OLS regression over the sample of the stalls selling hake that day. Each treatment assignment interacts with weekly dummies. We include strata fixed effects and cluster at the circuit level. The "No Enforcement" category is the omitted category, and it bunches observations assigned to the control group and the information campaign. To facilitate visual interpretation, we only present the confidence intervals associated with weeks one and four.

Tables

Table 3.1: Treatment Assignment

	No Enforcement	Enforcement	Total
No Information Campaign	9	41	50
Information Campaign	14	42	56
Total	23	83	106

	High Frequency Enforcement	Low Frequency Enforcement	Total
Predictable Enf. Schedule	19	20	39
Unpredictable Enf. Schedule	15	29	44
Total	34	49	83

Notes: The first panel shows the number of circuits assigned to each experimental cell jointly defined by the Information Campaign (row) and the Enforcement treatment (column). The second panel shows the number of observations in each enforcement sub-treatment.

Table 3.2: Treatment Effects on Hake Sales

VARIABLES	(1) Fresh, Visible Hake	(2) Any Hake Available (Hidden, Frozen, Visible)
Information Campaign Only	-0.133	-0.131
	(0.066)	(0.074)
Enforcement Only	-0.178	-0.130
	(0.082)	(0.089)
Info Campaign and Enforcement	-0.179	-0.139
	(0.074)	(0.094)
Change in Dep. Var. in Control Group During Intervention Period	-0.21	-0.36
N	901	901

Notes: This table reports the effect of each treatment arm on the availability of illegal hake fish. The variable Fresh Hake indicates when the hake was available fresh. Hake available indicates when was possible to buy fish in any form. The table reports marginal effects from a Probit regression. Other controls are included: municipality characteristics, strata fixed effects and the average level of the outcome variable in pre-intervention period. We control for pre-treatment values for the outcome variables in addition to the treatment indicator, because not all markets were visited in pre-intervention period. Robust standard errors clustered by circuit (the unit of randomization) in parentheses.

Table 3.3: Treatment Effects on Fish Consumption

VARIABLES	(1) Num. Times Hake Purchased	(2) Mention Ban (unprompted)
Information Campaign Only	-0.275	0.146
	(0.071)	(0.045)
Enforcement Only	-0.111	0.082
	(0.049)	(0.047)
Info Campaign and Enforcement	-0.098	0.107
	(0.046)	(0.051)
Mean Dep Var Control Group	0.49	0.07
N	3218	3319

Notes: This table presents the effect of different treatments on the reported consumption of hake fish during September 2015. The column 1 shows the marginal effects from a Poisson regression because the dependent variable is count data, the column 2 shows marginal effects from a Probit regression. Consumers were not asked about the ban, but surveyors registered if the ban was mentioned spontaneously. These regressions include socioeconomic characteristics and strata fixed effects. The numbers of observations in columns 1 and 2 differ because some consumers could not recall the number of times they purchased hake in the past month. Both Poisson and Probit are nonlinear models, and the average marginal effects of each treatment depend not only on the coefficients reported in this table, but also on the values of the covariates. Robust standard errors clustered by circuit in parentheses.

Table 3.4: Treatment Effect on Hake Sales by Enforcement Strategy

VARIABLES	(1)	(2)
	Any Hake Available (Fresh, Visible, Hidden or Frozen)	
Information Campaign only	-0.134	-0.135
	(0.073)	(0.072)
Enforcement on Predictable Schedule	-0.060	
	(0.083)	
Enforcement on Unpredictable Schedule	-0.192	
	(0.094)	
High Frequency Enforcement		-0.070
		(0.095)
Low Frequency Enforcement		-0.162
		(0.090)
p-value of Predictable = Unpredictable Sch.	0.036	
p-value of Low = High Int. Enf.		0.280
Change in Dep Var in Control		
During Intervention	-0.36	-0.36
N	901	901

Notes: This table presents the coefficient corresponding to the interaction term $T_c \times Post_t$ for each treatment. To retain statistical power, the cells "Enforcement only" and "Enforcement + Info Campaign" from Table 3.2 are combined under "Enforcement" and then sub-divided by schedule predictability (column 1), or intensity (column 2). So these coefficients should be interpreted as the average effects of enforcement when half the sample is also exposed to the information campaign. Note that we previously find evidence of null interaction effect between enforcement and info campaign (Muralidharan, Romero, and Wüthrich, 2019). Column 1 includes a dummy for the intensity sub-treatment, and column 2 includes a dummy for the predictability sub-treatment, but those coefficients are not shown. Each regression controls for the dependent variable in pre-intervention period, strata fixed effects and municipality characteristics. Probit regression marginal effects are reported. Robust standard errors clustered by circuit in parentheses..

Table 3.5: Treatment Effects on Hake Sales Controlling for Spillovers to Control Markets

VARIABLES	(1)	(2)	(3)	(4)
	Any Hake Available (Fresh, Visible, Hidden or Frozen)			
Enforcement on Predictable Schedule	-0.023	-0.030	-0.076	-0.058
	(0.083)	(0.069)	(0.080)	(0.060)
Enforcement on Unpredictable Schedule	-0.157	-0.167	-0.199	-0.177
	(0.091)	(0.075)	(0.084)	(0.084)
Spatial Spillover		-0.017		
(within 10 km of Treated market)		(0.082)		
Social Connection Spill-over			-0.071	
(Vendor knows a Treated Vendor)			(0.076)	
Supply-Chain Spill-over				-0.077
(Sources from same *Caleta* as Treated Vendor)				(0.081)
Change in Dep Var in Control During Intervention	-0.36	-0.36	-0.36	-0.36
N	901	901	901	901

Notes: This table re-estimates treatment effects controlling for possible spillover effects from treatment to control markets. We focus on enforcement treatments to ensure that the control cell size is large enough to be divided by exposure to spill-overs. We only present the coefficient corresponding to the interaction term $T_c \times Post_t$ for each treatment. Controls for $T_c, Post_t$, covariates, and baseline value of the dependent variable are included, but those coefficients are not shown. The table reports marginal effects from a Probit regression. The dependent variable is an indicator for any type of hake (fresh-visible, hidden or frozen) for sale in the stall. Robust standard errors are clustered by circuit, which was the unit of randomization. .

Table 3.6: Treatment Effect Transmission to Fishermen in Caletas

VARIABLES	(1) Earned Less in Sept 15 than Sept 14	(2) Feria Vendors buy less Hake in Sep15 compared to Sept 14	(3) Consumers are informed of Hake Ban
At least one circuit Enforced	0.238	0.169	-0.033
	(0.105)	(0.293)	(0.147)
Info Campaign	0.043	-0.101	0.343
	(0.158)	(0.322)	(0.186)
At least one circuit Enforced and Info Campaign	0.358	0.553	0.173
	(0.128)	(0.315)	(0.195)
Mean Dep Var Control Group	0.31	0.40	0.77
N	202	179	217

Notes: This table reports OLS coefficients based on fishermen responses. The variable Information campaign correspond to caletas located in municipalities assigned to receive any level of information campaign. The variable "At least one circuit enforced" considers all circuits located in the same municipality of the caleta. Socioeconomic variables of the caletas are included as covariates. In average, three fishermen were surveyed in each caleta. The numbers of observations in columns 1, 2 and 3 differ because some fishermen could not recall the earnings and vendor behavior in specific months. The dependent variables of each column are dummy variables. Robust standard errors clustered at caleta level in parentheses.

Table 3.7: Do Vendors Substitute to Selling Other Fish in Response to Treatment?

VARIABLES	(1) Pomfret Available	(2) Any Other Fish Available
Information Campaign Only	0.146	0.004
	(0.098)	(0.035)
Enforcement on Predictable Schedule	0.133	0.027
	(0.079)	(0.031)
Enforcement on Unpredictable Schedule	0.115	0.065
	(0.078)	(0.033)
Change in Dep Var in Control Markets During Intervention	0.29	0.09
N	901	6328

Notes: The table reports marginal effects from a Probit regression. The unit of observation in the first column is stall x secret shopper visit, and in the second column is stall × secret shopper visit × possible substitute fish variety. We only present the coefficient corresponding to the interaction term $T_c \times Post_t$ for each treatment. Controls for $T_c, Post_t$, covariates, and baseline value of the dependent variable are included, but those coefficients are not shown. Robust standard errors are clustered by circuit in parentheses.

Table 3.8: Treatment Effect on Fish Prices

VARIABLES	(1) Log Price Pomfret	(2) Log Price Substitute
Information Campaign Only	0.210	0.140
	(0.109)	(0.096)
Enforcement Only	-0.017	-0.021
	(0.066)	(0.055)
Info Campaign and Enforcement	0.081	0.047
	(0.065)	(0.059)
Change in Dep Var in Control During Intervention	-0.20	-0.27
N	614	939

Notes: The table reports treatment effects on hake substitutes' price from OLS regressions. The outcome variable is the log of price per kilo. The unit of observation in the first column is *stall with pomfret available × secret shopper visit*, and in the second column is *stall with any substitute available × secret shopper visit × substitute available fish variety*. We only present the coefficient corresponding to the interaction term $T_c \times Post_t$ for each treatment. Controls for $T_c, Post_t$, covariates, and baseline value of the dependent variable are included, but those coefficients are not shown. Robust standard errors are clustered by circuit in parentheses.

Table 3.9: Cost-Effectiveness Analysis

	(1) Reduction of Hake Sale	(2) Units of Hake Saved	(3) Implementation Costs (USD)	(4) Cost of Saving One Hake (USD)
Enforcement (*Overall*)	0.13	10,399	$ 62,900.25	$ 6.05
Unpredictable	0.192	15,358	$ 69,190.27	$ 4.51
Predictable	0.06	4,799	$ 62,900.25	$ 13.11
Low Frequency	0.162	12,959	$ 53,475.84	$ 4.13
High Frequency	0.07	5,599	$ 99,613.61	$ 17.79
Info Campaign	0.13	3,257	$ 16,213.53	$ 4.98

Notes: This table shows the benefits and costs of implementing each intervention. Column (1) reports the estimated effects (in percentage points) of treatments in the sale of any type of hake. Column (2) is computed based on the numbers of stall per feria, number of days a week the feria operate and number of fish available in a normal stall. Column (3) is reported by Sernapesca and represents a combination of fixed and variable costs. Finally, column (4) correspond to the ration of (3) over (2). These calculations assume the control group had zero enforcement nor information campaign. As we discussed in section C.3, the control group (mistakenly) received a few enforcement visits, the cost is negligible.

Bibliography

Philippe Aghion and Jean Tirole. Formal and real authority in organizations. *Journal of Political Economy*, 105(1):1–29, 1997. ISSN 00223808. doi: 10.1086/262063.

James Alm, Betty R Jackson, and Michael McKee. Getting the word out: Enforcement information dissemination and compliance behavior. *Journal of Public Economics*, 93(3): 392–402, 2009.

Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, 2001. ISSN 00129682. doi: 10.1111/1468-0262.00223.

Susan Athey, Jonathan Levin, and Enrique Seira. Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics*, 126(1):207–257, 2011. ISSN 00335533. doi: 10.1093/qje/qjq001.

Susan Athey, Dominic Coey, and Jonathan Levin. Set-asides and subsidies in auctions. *American Economic Journal: Microeconomics*, 5(1):1–27, 2013. ISSN 19457669. doi: 10.1257/mic.5.1.1.

Patrick Bajari. Comparing competition and collusion: A numerical approach. *Economic Theory*, 18(1):187–205, 2001. ISSN 09382259. doi: 10.1007/PL00004128.

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

Patrick Bajari, Robert McMillan, and Steven Tadelis. Auctions versus negotiations in procurement: An empirical analysis. *Journal of Law, Economics, and Organization*, 25(2): 372–399, 2009. ISSN 87566222. doi: 10.1093/jleo/ewn002.

Patrick Bajari, Stephanie Houghton, and Steve Tadelis. Bidding for Incomplete Contracts: An Empirical Analysis. *American Economic Review*, 104(4):1288–1319, 2014. doi: 10.3386/w12051. URL https://www.aeaweb.org/articles?id=10.1257/aer.104.4.1288.

Oriana Bandiera, Andrea Prat, and Tommaso Valletti. Active and passive waste in government spending. *American Economic Review*, 99(2009 Sept):1278–1308, 2009. ISSN

2038-2502 (Electronic). doi: 10.1257/aer.99.4.1278. URL http://econ.lse.ac.uk/staff/bandiera/bpv0208.pdf.

Abhijit Banerjee, Esther Duflo, Raghabendra Chattopadhyay, Daniel Keniston, and Nina Singh. The efficient deployment of police resources: Theory and new evidence from a randomized drunk driving crackdown in india. 2017.

Abhijit V Banerjee, Esther Duflo, and Rachel Glennerster. Putting a band-aid on a corpse: Incentives for nurses in the indian public health care system. *Journal of the European Economic Association*, 6(2-3):487–500, 2008.

Gary Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76:169–217, 1968.

Michael Carlos Best, Jonas Hjort, and David Szakonyi. Individuals and Organizations as Sources of State Effectiveness, and Consequences for Policy. 2017. doi: 10.3386/w23350. URL http://www.nber.org/papers/w23350.pdf.

Vivek Bhattacharya, James W. Roberts, and Andrew Sweeting. Regulating bidder participation in auctions. *The RAND Journal of Economics*, 45(4):675–704, 2014. doi: 10.1111/1756-2171.12067.

Christopher Blattman, Donald Green, Daniel Ortega, and Santiago Tobón. Pushing crime around the corner? estimating experimental impacts of large-scale security interventions. 2017.

Valentin Bolotnyy and Shoshana Vasserman. Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement. 2019.

William C Boning, John Guyton, Ronald H Hodge, Joel Slemrod, Ugo Troiano, et al. Heard it through the grapevine: Direct and network effects of a tax enforcement field experiment. 2018.

Erica Bosio, Simeon Djankov, Edward Glaeser, and Andrei Shleifer. Public Procurement in Law and Practice. 2020.

Bjarne Brendstrup and Harry J Paarsch. Nonparametric Estimation of Dutch and First-Price , Sealed-Bid Auction Models With Asymmetric Bidders. (319), 2003.

Jeremy Bulow and Paul Klemperer. Auctions Versus Negotiations. *American Economic Review*, 1996. ISSN 00028282. doi: 10.1126/science.151.3712.867-a.

Sebastian Calonico, Matias D. Cattaneo, and Rocio Titiunik. Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6):2295–2326, 2014. ISSN 1468-0262. doi: 10.3982/ecta11757.

Sandra Campo, Isabelle Perrigne, and Quang Vuong. Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics*, 18(2):179–207, 2003. ISSN 08837252. doi: 10.1002/jae.697.

Rodrigo Carril. Rules Versus Discretion in Public Procurement. 2019.

Paul Carrillo, Dina Pomeranz, and Monica Singhal. Dodging the taxman: Firm misreporting and limits to tax enforcement. *American Economic Journal: Applied Economics*, 9(2): 144–164, 2017.

Indranil Chakraborty, Fahad Khalil, and Jacques Lawarree. Competitive Procurement with Ex Post Moral Hazard. *RAND Journal of Economics*, 2020.

Victor Chernozhukov, Ivan Fernandez-Val, and Blaise Melly. Inference on Counterfactual Distributions. *Econometrica*, 81(6):2205–2268, 2013. ISSN 0012-9682. doi: 10.3982/ecta10582.

Raj Chetty, John Friedman, and Emmanuel Saez. Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings : Web Appendix Sample and Variable De ... nitions. *American Economic Review*, 103(7):2683–2721, 2013. ISSN 0002-8282. doi: 10.1257/aer.103.7.2683.

Raj Chetty, Mushfiq Mobarak, and Monica Singhal. Increasing tax compliance through social recognition. *Policy Brief*, 2014.

Pierre Andre Chiappori and Bernard Salanie. Testing Contract Theory: A Survey of Some Recent Work. 2002.

Christopher Costello, Daniel Ovando, Ray Hilborn, Steven D Gaines, Olivier Deschenes, and Sarah E Lester. Status and solutions for the world's unassessed fisheries. *Science*, 338 (6106):517–520, 2012.

Decio Coviello and Stefano Gagliarducci. Tenure in office and public procurement. *American Economic Journal: Economic Policy*, 9(3):59–105, 2017. ISSN 1945774X. doi: 10.1257/pol.20150426.

Decio Coviello and Mario Mariniello. Publicity requirements in public procurement: Evidence from a regression discontinuity design. *Journal of Public Economics*, 109:76–100, 2014. ISSN 00472727. doi: 10.1016/j.jpubeco.2013.10.008.

Francesco Decarolis. Awarding price, contract performance, and bids screening: Evidence from procurement auctions. *American Economic Journal: Applied Economics*, 6(1 A): 108–132, 2014. ISSN 19457782. doi: 10.1257/app.6.1.108.

Francesco Decarolis, Leonardo M. Giuffrida, Elisabetta Iossa, Vincenzo Mollisi, and Giancarlo Spagnolo. Bureaucratic Competence and Procurement Outcomes. *The Journal of Law, Economics, and Organization*, No. 24201, 2020. doi: 10.3386/w24201. URL http://www.nber.org/papers/w24201{%}0Ahttp://www.nber.org/papers/w24201.pdf.

Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from india. *The Quarterly Journal of Economics*, 128(4):1499–1545, 2013.

Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. The Value of Regulatory Discretion: Estimates From Environmental Inspections in India. *Econometrica*, 86 (6):2123–2160, 2018. ISSN 0012-9682. doi: 10.3982/ecta12876.

Florian Ederer, Richard Holden, and Margaret Meyer. Gaming and strategic opacity in incentive provision. *The RAND Journal of Economics*, 49(4):819–854, 2018.

Jan Eeckhout, Nicola Persico, and Petra E Todd. A theory of optimal random crackdowns. *American Economic Review*, 100(3):1104–35, 2010.

FAO. *Increasing the contribution of small-scale fisheries to poverty alleviation and food security*. Number 481. Food & Agriculture Organization of the United Nations., 2007.

FAO. Food and agriculture organization of the united nations, the state of world fisheries and aquaculture, 2014.

Drew Fudenberg and David K Levine. Learning in games and the interpretation of natural experiments. *American Economic Journal: Microeconomics (forthcoming)*, 2020.

François Gerard, Miikka Rokkanen, and Christoph Rothe. Bounds on treatment effects in regression discontinuity designs with a manipulated running variable. *Quantitative Economics*, 11(3):839–870, 2020. ISSN 1759-7323. doi: 10.3982/QE1079.

Edward L Glaeser and Andrei Shleifer. The rise of the regulatory state. *Journal of Economic Literature*, 41(2):401–425, 2003.

Victor P Goldberg. Competitive Bidding and the Production of Precontract Information. *The Bell Journal of Economics*, 8(1):250–261, 1977.

C. Gourieroux, A. Monfort, and E. Renault. Indirect inference. *Journal of Applied Econometrics*, 8(1 S):S85–S118, 1993. ISSN 10991255. doi: 10.1002/jae.3950080507.

Hugh Gravelle, Matt Sutton, and Ada Ma. Doctor behaviour under a pay for performance contract: treating, cheating and case finding?, 2010.

Wayne B Gray and Jay P Shimshack. The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence. *Review of Environmental Economics and Policy*, 5(1):3–24, 2011.

Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000. ISSN 00129682. doi: 10.1111/1468-0262.00123.

Raymond Guiteras, James Levinsohn, and Ahmed Mushfiq Mobarak. Encouraging sanitation investment in the developing world: A cluster-randomized trial. *Science*, 348(6237):903–906, 2015.

Philip A Haile and Yuichi Kitamura. Unobserved Heterogeneity in Auctions. *The Econometrics Journal*, 22(1):C1–C19, 2019. ISSN 1368-4221. doi: 10.1111/ectj.12121.

Jens Hainmueller, Michael J Hiscox, and Sandra Sequeira. Consumer demand for fair trade: Evidence from a multistore field experiment. *Review of Economics and Statistics*, 97(2):242–256, 2015.

Benjamin Hansen. Punishment and deterrence: Evidence from drunk driving. *American Economic Review*, 105(4):1581–1617, 2015.

Christopher Hansman, Jonas Hjort, and Gianmarco León. Interlinked firms and the consequences of piecemeal regulation. *Journal of the European Economic Association*, 2018.

Oliver Hart and John Moore. Incomplete Contracts and Renegotiation. *Econometrica*, 56(4):755–785, 1988.

Han Hong and Matthew Shum. Increasing Competition and the Winner ' s Curse : Evidence from Procurement. *Review of Economic Studies*, 69(4):871–898, 2002.

Hugo Jales and Zhengfei Yu. Identification and estimation using a density discontinuity approach. *Advances in Econometrics*, 38:29–72, 2017. ISSN 07319053. doi: 10.1108/S0731-905320170000038003.

Ginger Zhe Jin and Phillip Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451, 2003.

Matthew S Johnson. Regulation by shaming: Deterrence effects of publicizing violations of workplace safety and health laws. *American Economic Review (forthcoming)*, 2016.

Matthew S Johnson, David I Levine, and Michael W Toffel. Improving regulatory effectiveness through better targeting: Evidence from osha. 2019.

Karam Kang and Robert A Miller. Winning by Default: Why is there so Little COmpetition in Government Procurement? 2017.

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

Henrik Kleven. Bunching. *Annual Review of Economics*, 2016. doi: 10.1146/annurev-economics-080315-015234.

Henrik J Kleven and Mazhar Waseem. Using Notches To Uncover Optimization Frictions And Structural Elasticities: Theory and Evidence from Pakistan. *The Quarterly Journal of Economics*, 128(December):669–723, 2013. ISSN 00335533. doi: 10.1093/qje/qjt004. Advance.

Anja Kollmuss and Julian Agyeman. Mind the gap: Why do people act environmentally and what are the barriers to pro-environmental behavior? *Environmental Education Research*, 8(3):239–260, 2002.

Elena Krasnokutskaya. Identification and Estimation of Auction Models with Unobserved Heterogeneity. *Review of Economic Studies*, 78(1):293–327, 2011. ISSN 00346527. doi: 10.1093/restud/rdq004.

Elena Krasnokutskaya and Katja Seim. Bid Preference Programs and Participation in Highway Procurement Auctions. *American Economic Review*, 101(6):2653–2686, 2011. ISSN 00028282. doi: 10.1257/aer.101.6.2653.

Jean-Jacques Laffont and Jean Tirole. Adverse Selection and Renegotiation in Procurement. *The Review of Economic Studies*, 57(4):597, 1990. ISSN 00346527. doi: 10.2307/2298088.

Jean-Jacques Laffont and Jean Tirole. *A Theory of Incentives in Procurement and Regulation.* MIT press, 1993.

Edward P Lazear. Speeding, terrorism, and teaching to the test. *The Quarterly Journal of Economics*, 121(3):1029–1061, 2006.

Bernard Lebrun. First price auctions in the asymmetric N bidder case. *International Economic Review*, 40(1):125–142, 1999. ISSN 00206598. doi: 10.1111/1468-2354.00008.

Dan Levin and James L Smith. Equilibrium in Auctions with Entry. *American Economic Review*, 84(3):585–599, 1994.

Gregory Lewis and Patrick Bajari. Procurement Contracting with Time Incentives Theory and Evidence. *Quarterly Journal of Economics*, 2011. doi: 10.1093/qje/qjr026.Advance. URL http://www.jstor.org/stable/10.2307/137292.

Tong Li and Xiaoyong Zheng. Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions. *Review of Economic Studies*, 76(4):1397–1429, 2009. ISSN 00346527. doi: 10.1111/j.1467-937X.2009.00558.x.

Tong Li and Xiaoyong Zheng. Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions. *Journal of Econometrics*, 168(1): 29–46, 2012. ISSN 03044076. doi: 10.1016/j.jeconom.2011.09.004. URL http://dx.doi.org/10.1016/j.jeconom.2011.09.004.

Jeffrey B. Liebman and Neale Mahoney. Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement Faculty Research Working Paper Series. *American Economic Review*, 107(11):3510–3549, 2017. doi: 10.1145/2462456.2464460.

Robert E Lucas. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. Elsevier, 1976.

Alexander Mackay. Contract Duration and the Costs of Market Transactions. 2018.

W Bentley Macleod and James M Malcomson. Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. *Econometrica*, 1989.

James M. Malcomson. Relational incentive contracts. *The Handbook of Organizational Economics*, (508):1014–1065, 2012. doi: 10.1515/9781400845354-027.

Vadim Marmer, Artyom Shneyerov, and Pai Xu. What model for entry in first-price auctions? A nonparametric approach. *Journal of Econometrics*, 176(1):46–58, 2013. ISSN 03044076. doi: 10.1016/j.jeconom.2013.04.005. URL http://dx.doi.org/10.1016/j.jeconom.2013.04.005.

Eric Maskin and John Riley. Asymmetric Auctions. *Review of Economic Studies*, 67(3): 413–438, 2003a. ISSN 0034-6527. doi: 10.1111/1467-937x.00137.

Eric Maskin and John Riley. Equilibrium in Sealed High Bid Auctions. *Review of Economic Studies*, 67(3):439–454, 2003b. ISSN 0034-6527. doi: 10.1111/1467-937x.00138.

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

Edward Miguel and Michael Kremer. Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217, 2004.

Karthik Muralidharan, Mauricio Romero, and Kaspar Wüthrich. Factorial designs, model selection, and (incorrect) inference in randomized experiments. 2019.

Joana Naritomi. Consumers as tax auditors. *CEPR Discussion Paper DP13276*, 2018.

Deniz Okat. Deterring fraud by looking away. *The RAND Journal of Economics*, 47(3): 734–747, 2016.

Paul Oyer. Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality. *The Quarterly Journal of Economics*, 113(1):149–185, 1998.

Ariel Pakes and David Pollard. Simulation and the Asymptotics of Optimization Estimators. *Econometrica*, 57(5):1027–1057, 1989.

Zhuan Pei and Yi Shen. The devil is in the tails: Regression discontinuity design with measurement error in the assignment variable. *Advances in Econometrics*, 38(October): 455–502, 2017. ISSN 07319053. doi: 10.1108/S0731-905320170000038019.

Dina Pomeranz. No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, 105(8):2539–69, 2015.

Robert H Porter and J Douglas Zona. Detection of Bid Rigging in Procurement Auctions. *Journal of Political Economy*, 101(3):518–538, 1993.

Ritva Reinikka and Jakob Svensson. Fighting corruption to improve schooling: Evidence from a newspaper campaign in uganda. *Journal of the European Economic Association*, 3 (2-3):259–267, 2005.

Nicholas Ryan. Contract Enforcement and Productive Efficiency: Evidence From the Bidding and Renegotiation of Power Contracts in India. *Econometrica*, 88(2):383–424, 2020. ISSN 0012-9682. doi: 10.3982/ecta17041.

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

William F Samuelson. Competitive Bidding with Entry Costs. *Economics Letters*, 17:53–57, 1985.

Abebe Shimeles, Daniel Zerfu Gurara, and Firew Woldeyes. Taxman's dilemma: Coercion or persuasion? evidence from a randomized field experiment in ethiopia. *American Economic Review*, 107(5):420–24, 2017.

Jay P Shimshack and Michael B Ward. Regulator reputation, enforcement, and environmental compliance. *Journal of Environmental Economics and Management*, 50(3):519–540, 2005.

Larissa Roxanna Smith and Víctor M Muñiz-Fraticelli. Strategic shortcomings of the dodd-frank act. *The Antitrust Bulletin*, 58(4):617–633, 2013.

Daniel F Spulber. Auctions and Contract Enforcement. *The Journal of Law, Economics, and Organization*, 6(2), 1990. ISSN 8756-6222. doi: 10.1093/oxfordjournals.jleo.a036995.

Robert N. Stavins. The problem of the commons: Still unsettled after 100 years. *American Economic Review*, 101(1):81–108, 2011.

Rainer Storn and Kenneth Price. Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*, pages 341–359, 1997.

Subpesca. Propuesta de política pública de desarrollo productivo para la pesca artesanal. *Servicio Nacional de Pesca de Chile*, 2013.

Subpesca. Estado de situación principales pesquerías chilenas. *Servicio Nacional de Pesca de Chile*, 2015.

Andrew Sweeting and Vivek Bhattacharya. Selective Entry and Auction Design. *International Journal of Industrial Organization*, 43:189–207, 2015. ISSN 01677187. doi: 10.1016/j.ijindorg.2015.03.004. URL http://dx.doi.org/10.1016/j.ijindorg.2015.03.004.

Ferenc Szucs. Discretion and Corruption in Public Procurement. 2020.

Jean Tirole. Incomplete Contracts: Where do we Stand? *Econometrica*, 67(4):741–781, 1999.

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

WWF. Estimación de la pesca inn en la pesquería de merluza común. *World Wide Fund for Nature*, 2017.

Appendix A

Appendix: Competition under Incomplete Contracts and the Design of Procurement Policies I: Effects of Publicity

A.1 Additional Figures

Figure A.1.1: Contract Timeline and Data Sources

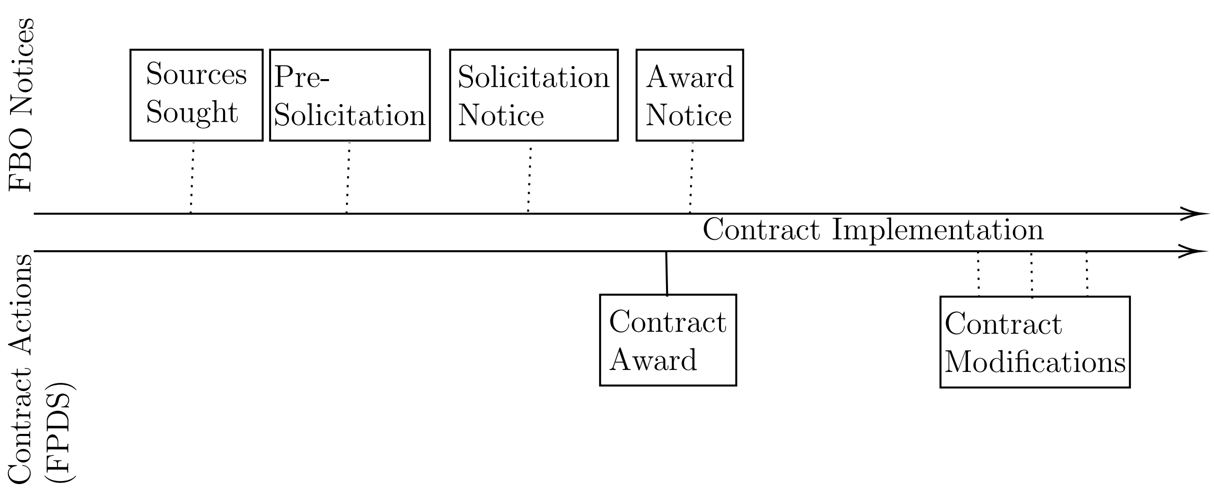

Notes: This figure presents a timeline of events associated with a typical contract. Milestones located above the arrows correspond to notices that are published on the government's point of entry (`fedbizopps.gov`). Milestones below the arrows generate information that is recorded on the Federal Procurement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

Figure A.1.2: Intuition of Method to Estimate Mean Price Effects

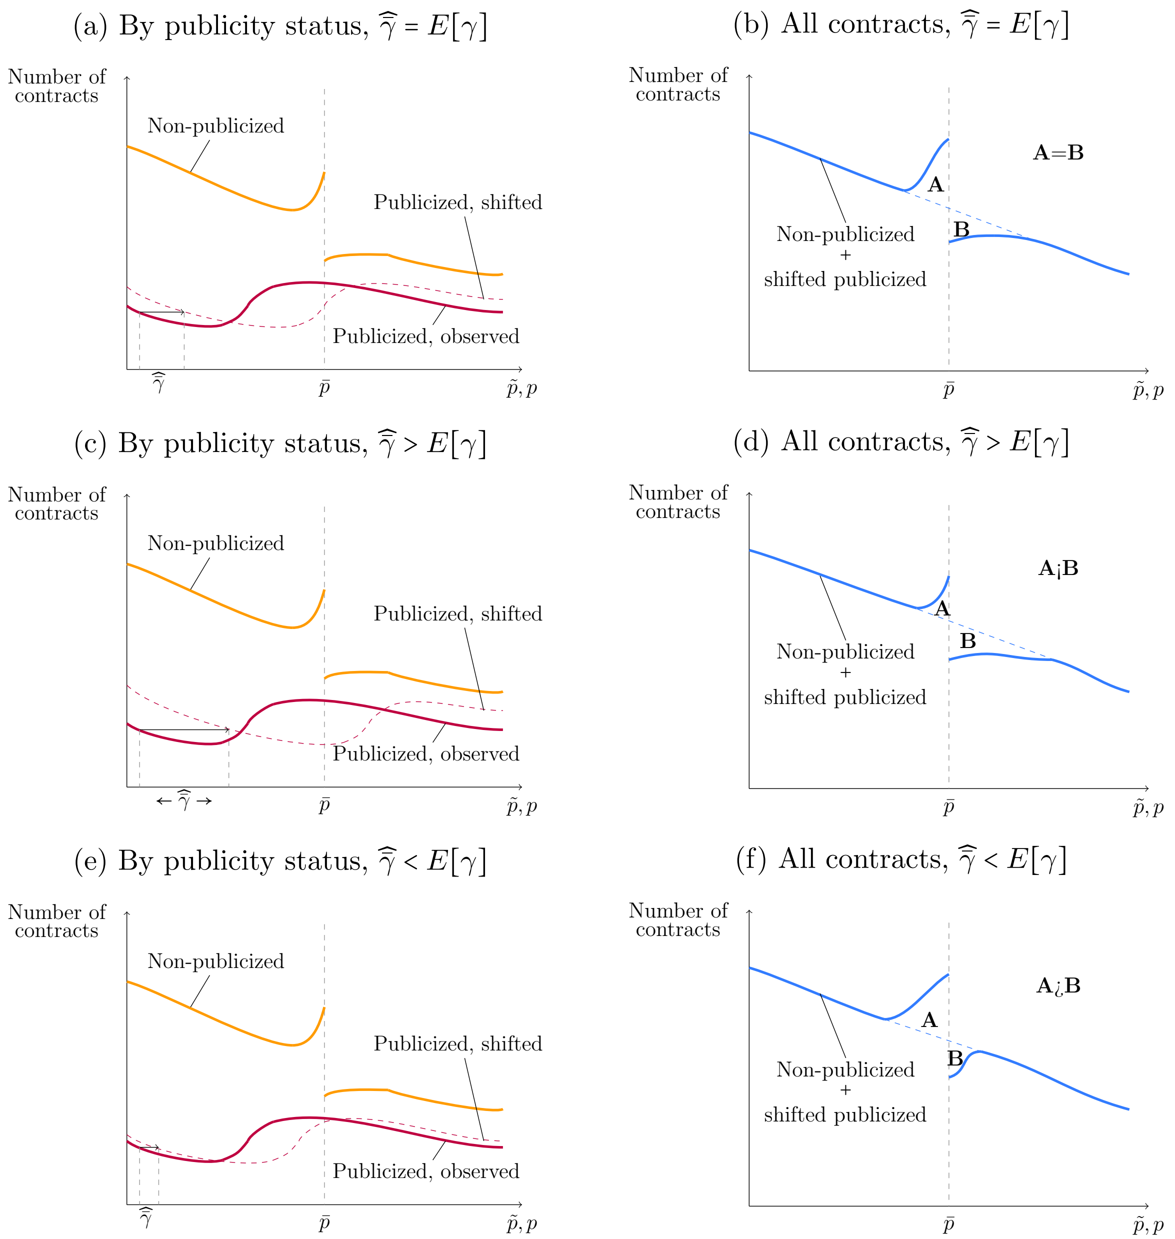

Notes: This figure provides (graphical) intuition of the procedure to estimate the mean price effect based on the integration constraint condition, i.e., the sum of excess of mass below the threshold equals the sum of missing masses above the threshold. Panels (a), (c), and (e) display distributions of publicized and non-publicized contracts. Panels (b), (d), and (f) show the corresponding overall distributions, i.e., the blue line in panel (b) corresponds to the sum of the yellow and red lines in panel (a). The key intuition is that the integration constraint condition is only met if the distribution of publicized contracts is re-centered by the correct mean of price effect, i.e., the resulting distribution has mean zero.

Figure A.1.3: Intuition of Method to Estimate Ex-Ante Price Distributions

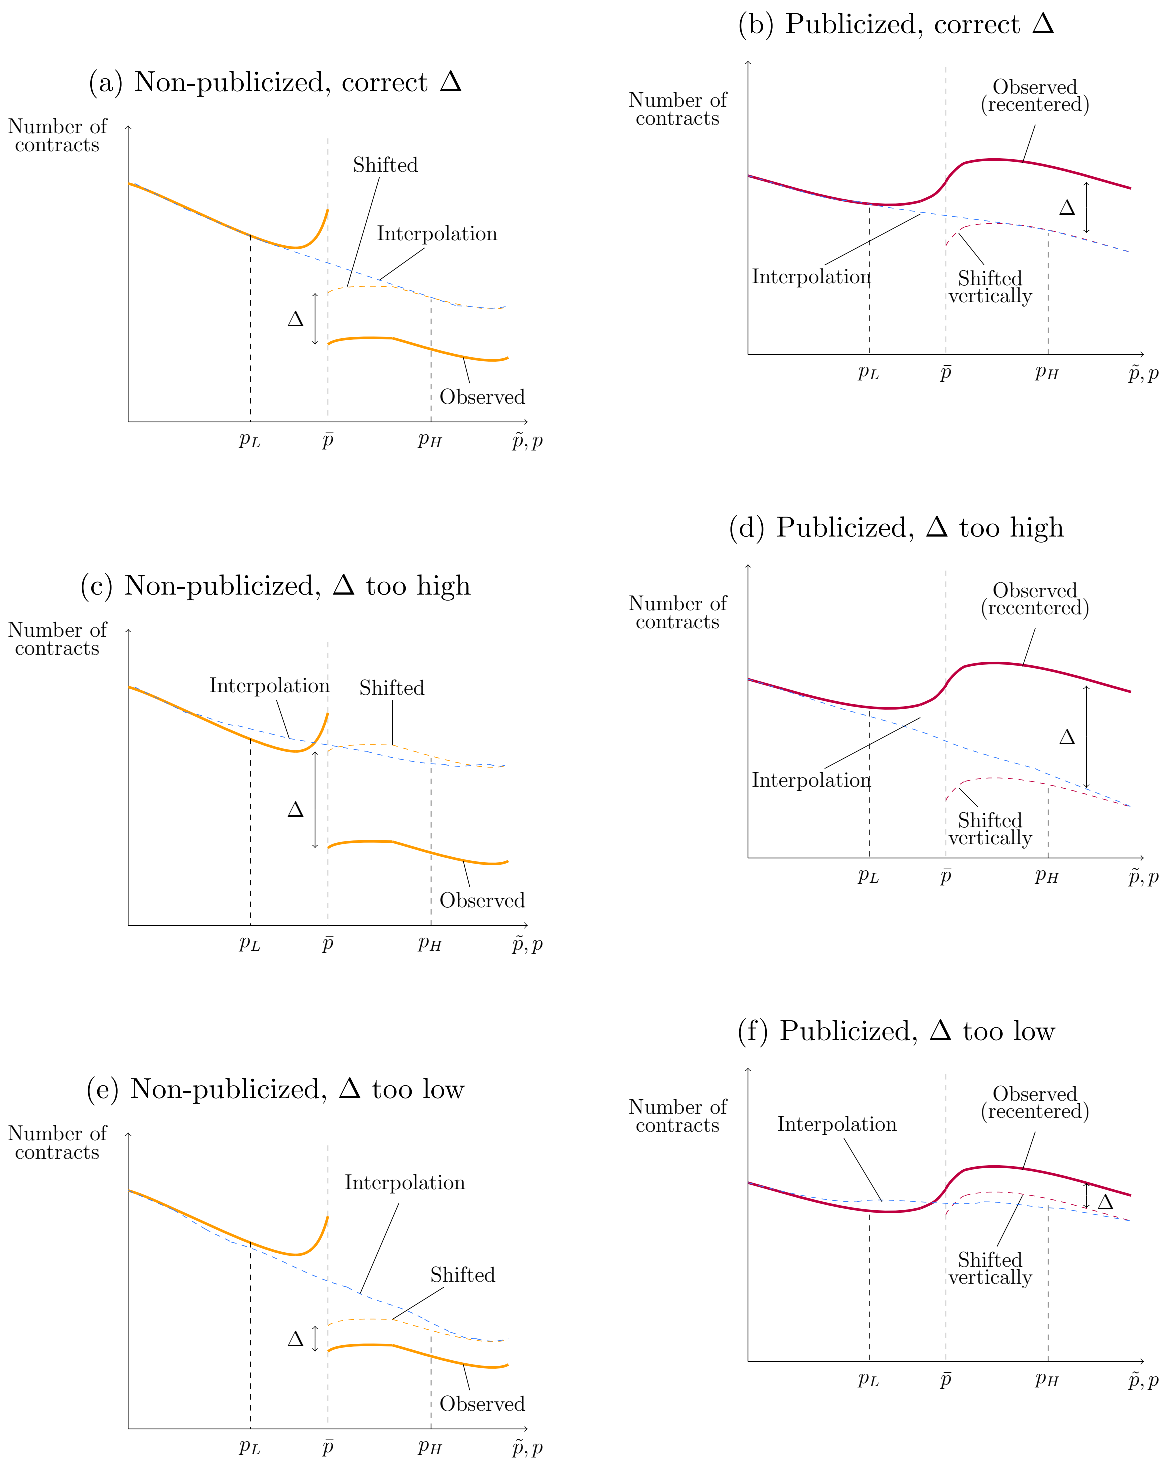

Notes: This figure provides (graphical) intuition of the procedure to estimate the ex-ante price distribution. The method considers identifying the discrete change in the distribution of publicized contracts (Δ) that matches with the drop in the distribution of non-publicized contracts. Panels (a), (c), and (e) display distributions of non-publicized contracts. Panels (b), (d), and (f) show the distributions of publicized contracts. The procedure builds upon the general interpolation (dashed blue line) that relates the distributions of publicized and non-publicized contracts. We recover the Δ by identifying the vertical shift of the distributions that matches the

Figure A.1.4: Pre-award characteristics around the threshold

(a) Last month Fiscal Year

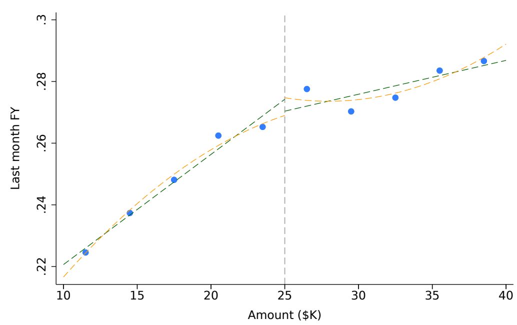

(b) Set-asides

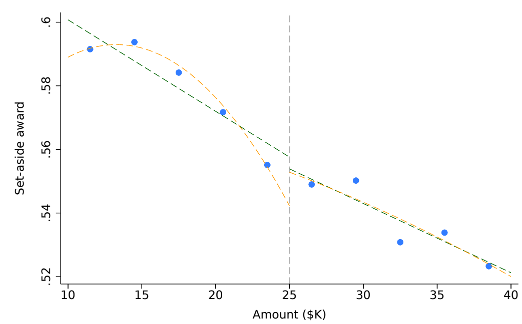

(c) Simplified acquisition procedures

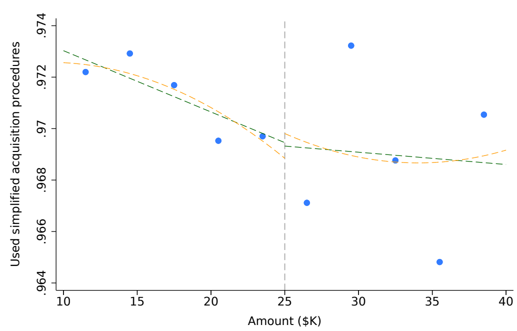

(d) Service contracts

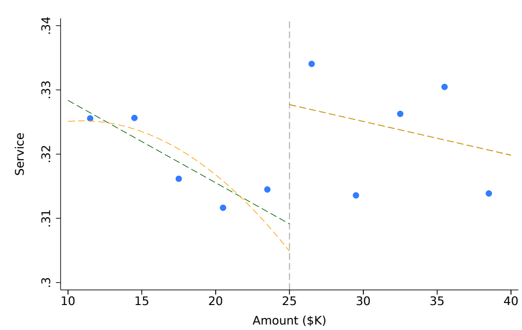

Notes: This figure presents four binned scatter plots, which depict an average pre-award characteristic by bins of award amounts, as well as linear and quadratic fits at each side of \$25,000. The pre-award characteristic in each Panel is as follows: (a) an indicator equal to one if the contract was solicited the last month of the fiscal year (September); (b) an indicator equal to one if the contract was set-aside for a preferential group (e.g. small businesses); (c) an indicator equal to one if the contract was awarded using simplified acquisition procedures; (d) an indicator equal to one if the award is for a service contract. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between \$ 10,000 and \$ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of \$3,000 dollars length.

Figure A.1.5: Publicity Effects on Post-Award Contract Performance

(a) Delays (days)

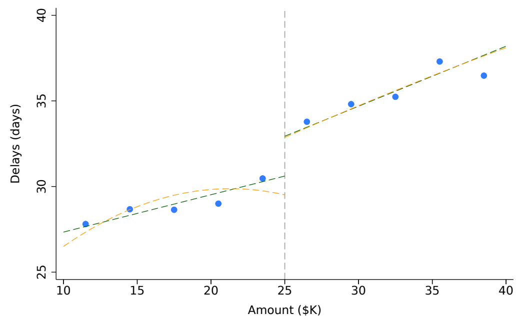

(b) Cost-Overruns (Share of Award Value)

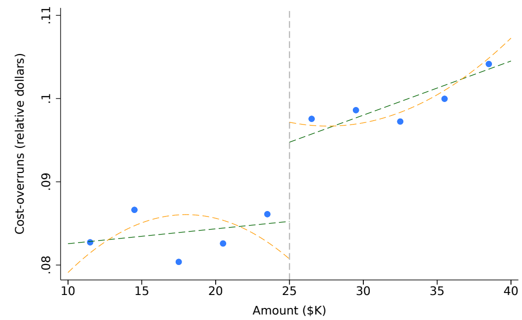

(c) Number of Contract Modifications (Ex-Post)

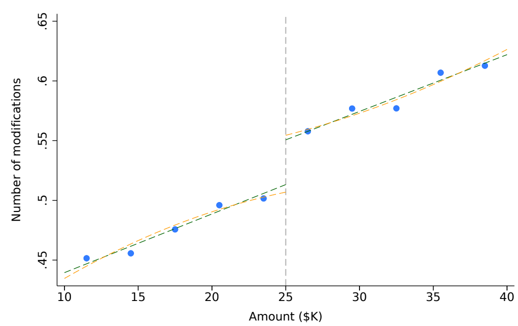

Notes: This figure presents four binned scatter plots, which depict an average post-award characteristic by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The pre-award characteristic in each Panel is as follows: (a) number of days of contract implementation delays; (b) cost-overruns as a share of award value; (c) number of modification to the original contract. The data source is the Federal Procurement Data System-Next Generation. The sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure A.1.6: Heterogeneous publicity adoption by major departments

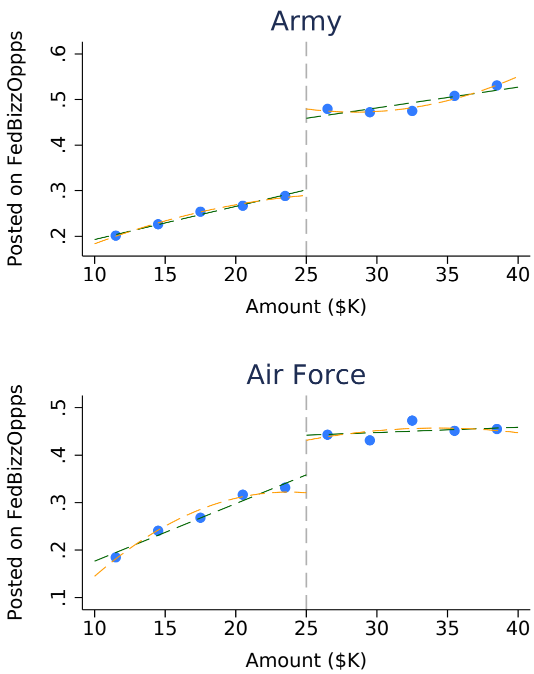

Notes: This figure presents three binned scatter plots, which depict the share of contracts publicized in FedBizzOpps by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists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

Figure A.1.7: Heterogeneous effects on competition by major departments

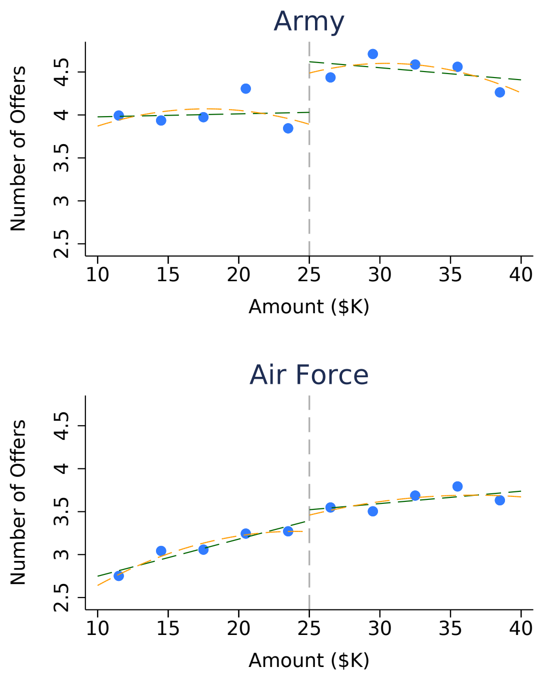

Notes: This figure presents three binned scatter plots, which depict the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D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

Figure A.1.8: Heterogeneous effects on winner characteristics by major departments

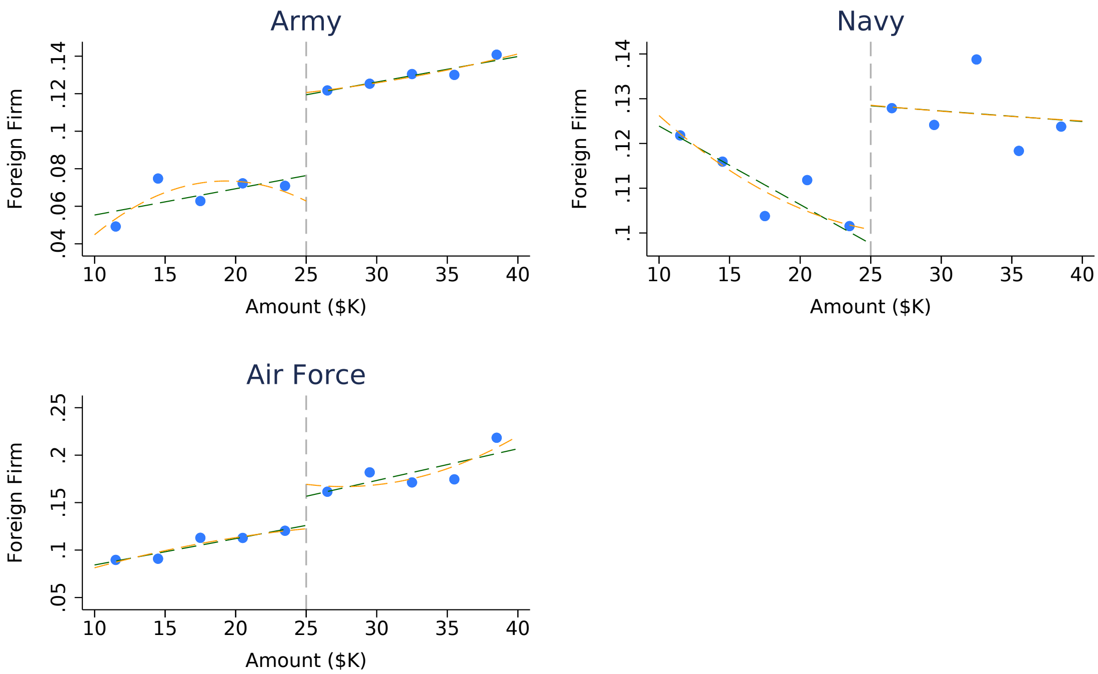

Notes: This figure presents three binned scatter plots, which depict the share of contracts awarded to a foreign firm by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists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

Figure A.1.9: Heterogeneous effects on performance by major departments

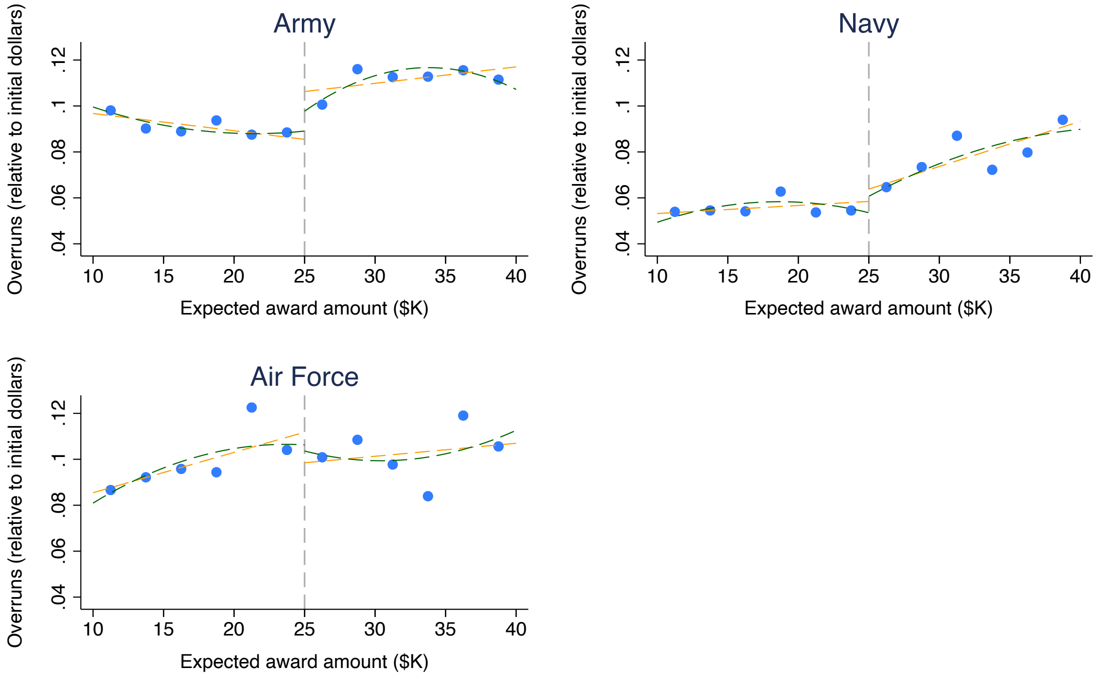

Notes: This figure presents three binned scatter plots, which depict average cost overruns by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the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

Figure A.1.10: Heterogeneous publicity adoption: goods versus services

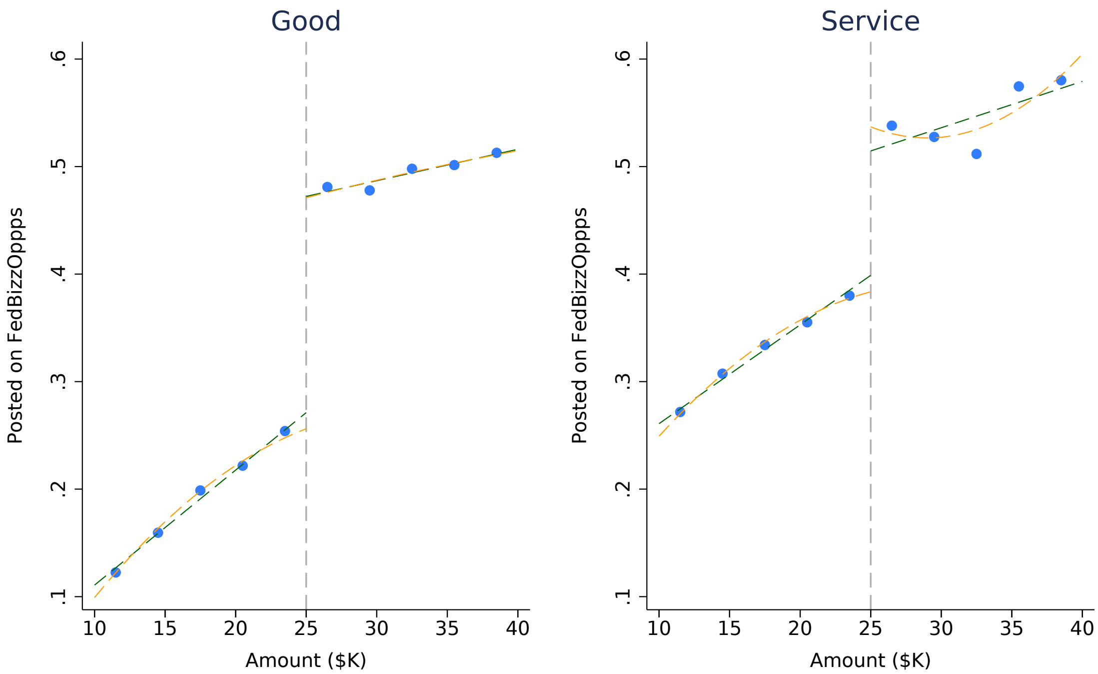

Notes: This figure presents two binned scatter plots, which depict the share of publicized contracts by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure A.1.11: Heterogeneous effects on competition: goods versus services

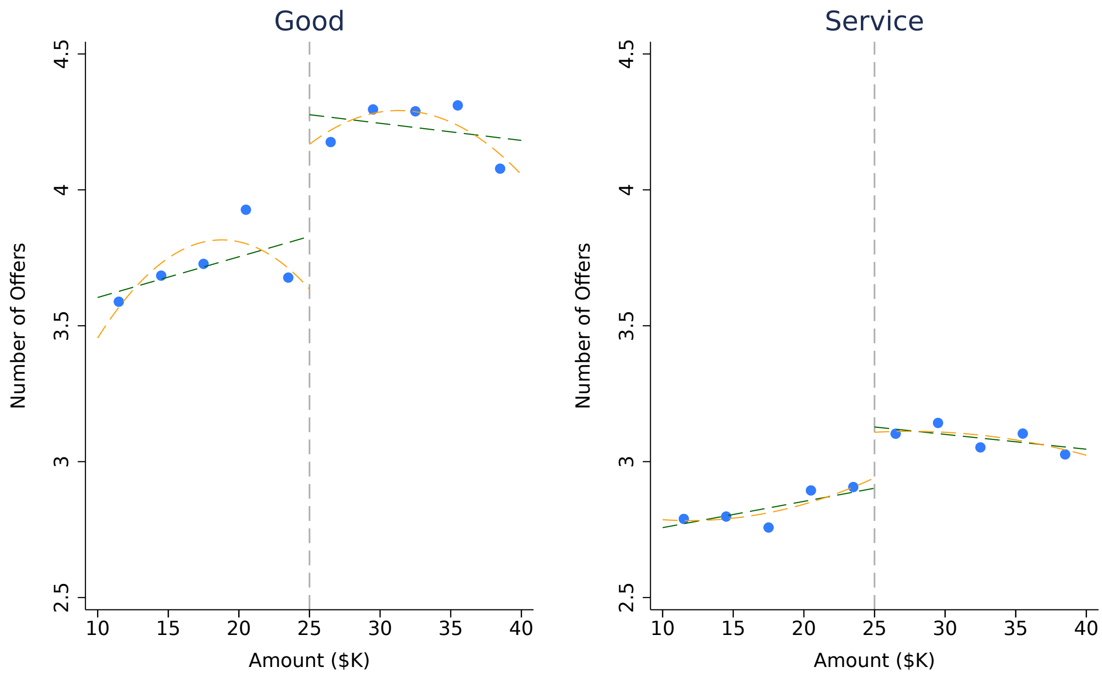

Notes: This figure presents two binned scatter plots, which depict the average number of offers received by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

Figure A.1.12: Heterogeneous effects on winner characteristics: goods versus services

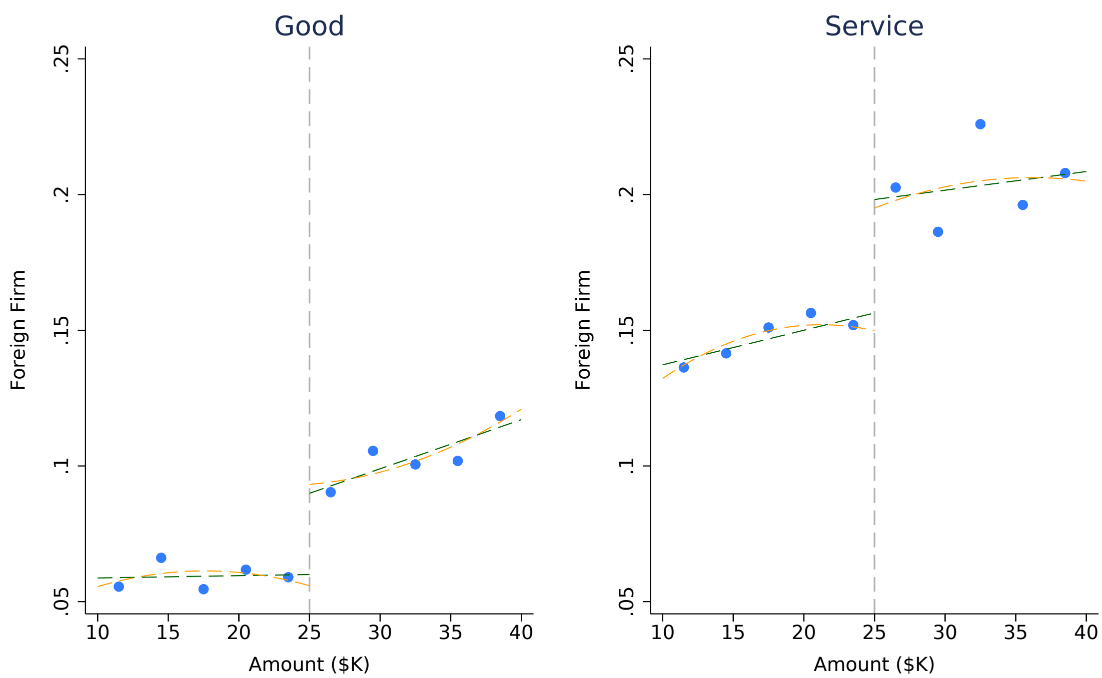

Notes: This figure presents two binned scatter plots, which depict the share of contracts awarded to a foreign firm by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $2,500 dollars length.

Figure A.1.13: Heterogeneous effects on performance: goods versus services

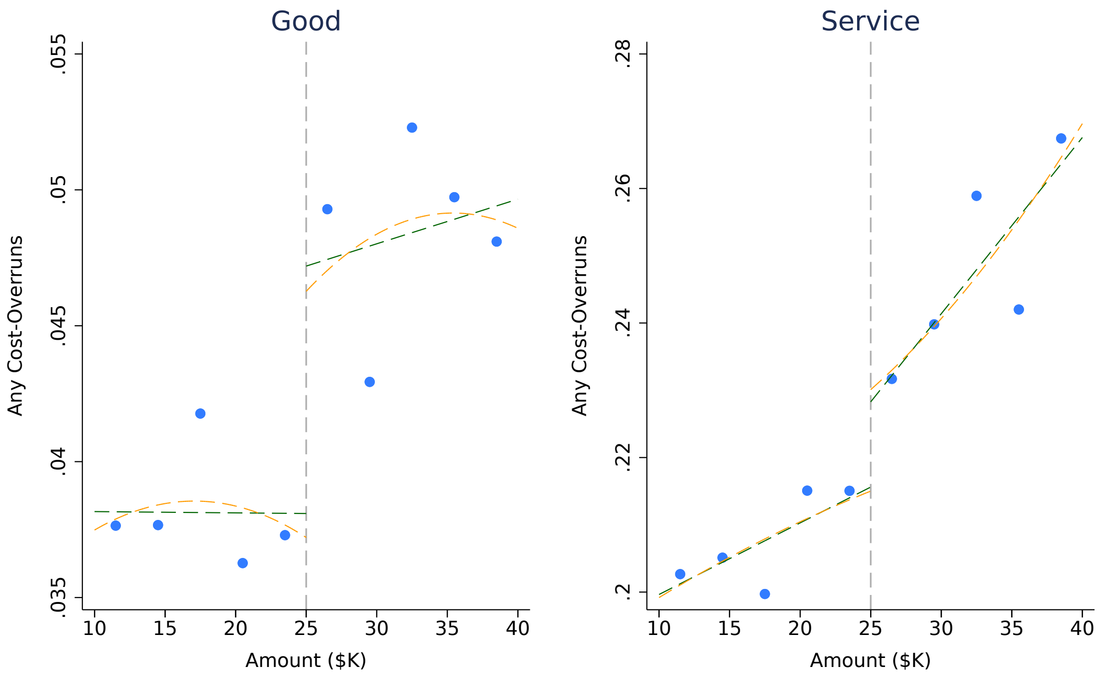

Notes: This figure presents two binned scatter plots, which depict share of contracts with cost overruns by bins of award amounts, as well as linear and quadratic fits at each side of $25,000. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the time of the award, divided by expected obligations. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between $ 10,000 and $ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019. Panel (a) restricts the sample to awards for goods, while Panel (b) restricts the sample to service contracts. Award amounts are discretized into right-inclusive bins of $3,000 dollars length.

A.2 Additional Tables

Table A.2.1: Effect on Cost-Overruns Controlling for Firm-FE

	(1)	(2)	(3)	(4)
	Share with Cost Overruns			
Estimate	0.076	0.099	0.022	0.044
S. E.	(0.027)	(0.042)	(0.029)	(0.035)
Estimation Method	IV	CCT	IV	CCT
Firm Fixed Effect Included	No	No	Yes	Yes
Number of Observations	69296	69296	69296	69296

Notes: This table shows the instrumental variable estimates for appearing in FedBizzOpps. The dependent variable is an indicator of having any positive cost-overrun. Columns 1 and 3 are linear IV estimates. Columns 2 and 4 show fuzzy RD estimates using robust-local polynomial regression (CCT - Cataneo et al. 2014). Columns 3 and 4 include firm fixed effects to control for the average performance of the contractors. The sample consists of observations from contractors that appear more than once in the data; otherwise, they are naturally dropped from fixed effect regression. Column 4 is estimated using a residualized outcome variable. Cost overruns are computed as the difference between actual obligated contract dollars and expected total obligations at the award time. The data source is the Federal Procurement Data System-Next Generation. The full sample consists of non-R&D definitive contracts and purchase orders, with award values between \$ 10,000 and \$ 40,000, awarded by the Department of Defense in fiscal years 2015 through 2019.

A.3 Additional Details on the Setting

FedBizzOpps

FedBizOpps.gov (FBO) has been designed as a single government point of entry (GPE) for Federal buyers to publish and for vendors to find posted Federal business opportunities across departments and agencies. The FAR (part 5) regulates the publicity of contract actions. The goals of publicity policy (FAR 5.002) are (a) increase competition, (b) broaden industry participation in meeting Govt requirements (c) assist small businesses (and VO, VOSD, WO, HUBZone, etc.) in winning contracts and subcontracts. The FAR requires that contract actions expected to exceed \$25,000 must be *synopsized* in the GPE. Contract actions under \$25,000 must publicize "by displaying in a public place, or by any appropriate electronic means." The contracting officer is exempted to advertise in GPE (FAR 5.102(a)5 and 5.202), when "disclosure compromises national security, " "nature of the file (e.g., size) does not make it cost-effective or practicable," the "agency's senior procurement executive makes a written determination that it is not in the Government's interest," and several other special cases (see FAR 5.202).

Figure A.3.1 displays screenshots to the website. Panel (a) shows the list of opportunities, Panel (b) includes the information contained a specific solicitation :

Figure A.3.1: FedBizzOpps

(a) List of Opportunities (b) Example Solicitation

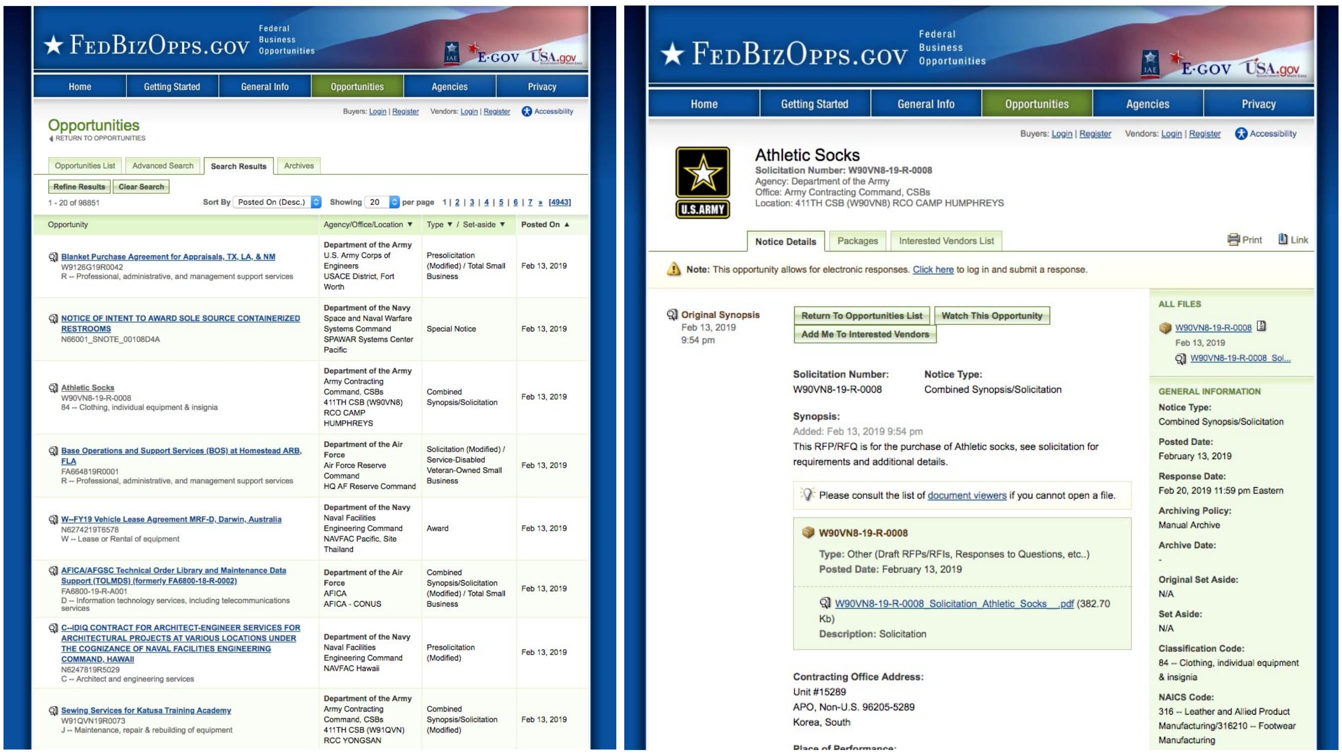

Notes: These figures show screenshots to FBO.gov. Panel (a) the list of opportunities and search alternatives. Panel (b) shows a particular solicitation for athletic socks, required by an Army procurement office. These screenshots were captured on Feb 13, 2019.

Types of FBO Notices

There are two broad types of FBO notices: *pre-award* and *post-award* notices. The *pre-award* notices are divided into four actions:[1]

- **Presolicitation:** The pre-solicitation notice makes vendors aware that a solicitation may follow. Vendors may add themselves to the Interested Vendors List, if the posting agency has enabled this feature. This helps government agencies determine if there are

[1]Here we omit uncommonly used actions: *Sale of Surplus Property*, *Justification and Approval (J&A)*, *Fair Opportunity / Limited Sources Justification*, *Foreign Government Standard*, and *Intent to Bundle Requirements (DoD-Funded)*.

qualified vendors to perform the work scope and allows the contracting office to gather information on the interested vendors.

- **Combined Synopsis/Solicitation:** Most opportunities classified this way are open for bids from eligible vendors. These opportunities include specifications for the product or service requested and a due date for the proposal. The notice will specify bidding procedures in the details of the solicitation.
- **Sources Sought:** The Sources Sought notice is a synopsis posted by a government agency seeking possible sources for a project. It is not a solicitation for work or a request for proposal. For more information, see FAR 7.3 and OMB Circular A-76.
- **Special Notice:** Agencies use Special Notices to announce events like business fairs, long-range procurement estimates, pre-bid/pre-proposal conferences, meetings, and the availability of draft solicitations or draft specifications for review.

The post-award notices are essentially *award notices*:

- **Award Notice:** When a federal agency awards a contract in response to a solicitation, they may choose to upload a notice of the award to allow the interested vendors to view the vendor receiving the awarded contract, and amount agreed upon.

Figure A.1.1 describes the life-cycle of a project and how different stages are linked to FBO actions.

Dataset Details

Our analysis combines data from two sources: Federal Procurement Data System - Next Generation (FPDS-NG) and data scrapped directly from FedBizzOpps.gov (FBO).
FPDS-NG. The FPDS-NG tracks the universe of federal awards that exceed \$5,000.[2] The Federal Acquisition Regulation (FAR) requires Contracting Officers (COs) must submit complete reports on all contract actions. Thus, every observation corresponds to a contract action, representing either an initial award or a follow-on action, e.g., modification, termination, renewal, or exercise of options. For each observation, we observe detailed information, such as the dollar value of the funds obligated by the transaction; a four-digit product category code (PSC); six-digit Industry (NAICS) code; identification codes for the agency, sub-agency, and contracting office making the purchase; the identity of the private vendor (DUNS); the type of contract pricing (typically, fixed-price or cost-plus); the extent of competition for the award; characteristics of the solicitation procedure; the number of offers received; and the applicability of a variety of laws and statutes. We collapse all actions by contract ID. As a reference, 80% of awarded contracts are smaller than $\$50,000$.

[2] The data can be downloaded from `usaspending.gov`

Our analysis contemplates overruns in terms of cost and time of completion. We define contract delays and cost overruns based on related literature (Decarolis et al., 2020). We exclude outliers on both variables as they are likely associated with data entry issues. We cross-checked dates and amounts for contract award notices that appeared in FBO and found that mismatches are uncommon.

FBO Data. We use daily archives of all information posted in FBO. Every data row corresponds to a different notice action. Each action is associated with a unique URL. The two primary IDs to match FBO data with other datasets are "solicitation number" and "contract award number. The former identifies pre-award actions, whereas award notices are identified using "contract award number." A relevant fraction of the award-notices are not linked with any of the pre-award notices. FPDS data contain both IDs. Roughly, an annual database contains 300,000 notices.

The data preparation consists in three steps; first, we clean IDs and classify different actions associated with each ID. Second, we merge with FPDS data using contract number, then update solicitation number when both exist, finally merge and append unmatched observations using solicitation number. The last step is to collapse the data at the FPDS contract ID level. So the resulting dataset contains all the contract ids that also appeared in FBO.

We define that a contract appeared in FBO (treatment indicator) if the contract award has a solicitation number associated with at least one of the FBO pre-award actions described above.

A.4 Empirical Framework for Estimating the Effects of Publicity on Contract Outcomes

This Appendix presents a detailed exposition of the empirical framework introduced in section 1.3. Section A.4 presents our theoretical framework and the set of results that motivate the density analysis. Section A.4 explains the density analysis in detail, including all implementation details. Section A.4 discusses how to correct naive RDD estimates to account for price effects and potential measurement error. Section A.4 explains how we account for potential bunching responses in the RDD framework.

Model

Preliminaries

Consider a series of observed contract awards $t \in \{1, .., T\}$. Let $\tilde{p}_t$ be the *ex-ante award price* of contract t, which corresponds to the agency's estimate of what the contract price will be. Let p_t be the *observed award price* of contract t. $\tilde{p}_t$ and p_t are normalized relative to a policy threshold of \$25,000 and measured in logs. Therefore, negative (positive) values of $\tilde{p}_t$ and p_t are said to be below (above) the threshold for the purpose of the policy described below.

Prior to the award, the buyer decides whether to publicize the solicitation ($D_t = 1$) or not ($D_t = 0$). Let $p_t^d(\tilde{p}_t)$ be the potential price that we would observe for contract t, given an ex-ante estimate of $\tilde{p}_t$ and a publicity decision $D_t = d$, for $d \in \{0, 1\}$. There is a policy that encourages buyers to choose $D_t = 1$ for awards expected to exceed the threshold (i.e. for $\tilde{p}_t > 0$).

The buyer may choose to *strategically bunch* ($B_t = 1$), which means that she modifies the characteristics of the initial purchase, in order to obtain an award price equal to $p_t^B(\tilde{p}_t)$, choosing $D_t = 0$ without being affected by the policy. $p_t^B(\tilde{p}_t)$ is equal to, or slightly below 0.

Therefore, observed prices can be written as:

$$p_t = p_t^0(\tilde{p}_t) + D_t \cdot \left[p_t^1(\tilde{p}_t) - p_t^0(\tilde{p}_t)\right] + B_t \cdot (1 - D_t) \cdot \left[p_t^B(\tilde{p}_t) - p_t^0(\tilde{p}_t)\right]$$

We assume the following:

A1 $\tilde{p}_t$ are i.i.d. draws from a distribution with smooth density $f_{\tilde{p}}(\cdot)$.

A2 $p_t^0(\tilde{p}_t) = \tilde{p}_t + \xi_t$, with $\xi_t \sim F_\xi(\cdot)$, $E[\xi_t] = 0$, and $\xi_t \perp \tilde{p}_t$.

A3 $p_t^1(\tilde{p}_t) = \tilde{p}_t + \gamma_t$, with $\gamma_t \sim F_\gamma(\cdot)$, $\gamma_t \perp \tilde{p}_t$, and $\gamma_t \perp \xi_t$.

A4 $\Pr(D_t = 1|\tilde{p}_t) \equiv \tilde{\pi}_D(\tilde{p}_t) = \tilde{\pi}_D^*(\tilde{p}_t) + \delta \cdot \mathbf{1}[\tilde{p}_t > 0]$, for a continuous function $\tilde{\pi}_D^*(\cdot)$.

A5 There exist $p_H > 0$ such that $B_t = 0$ for all $\tilde{p}_t > p_H$.

Note that here we present a slightly more general version of the model that in Section 1.3. In particular, **A2** allows for measurement error in agencies' ex-ante estimates.

Discretizing award values

Consider the division of the range of possible (normalized) award values into a set of equally-sized and right-inclusive bins around the threshold $b \in \{-R, (-R+1), ..., -1, 0, 1, ..., (R-1), R\}$. Note that bin $b = 0$ includes awards right at, or slightly below, the policy threshold.

Let $\{n_b^d\}_{b=-R}^R$ be the frequency distribution of observed awards conditional on treatment (publicity) status $D_t = d$, for $d \in \{0, 1\}$, so that n_b^d denotes the number of contracts with treatment status d and observed award value $p_t \in b$. Likewise, let $\{\tilde{n}_b^d\}_{b=-R}^R$ represent the (unobserved) frequency distribution of latent ex-ante prices. We also denote the distribution of *all* awards (both publicized and non-publicized) by simply omitting the superscript. That is, $n_b = n_b^0 + n_b^1$, and $\tilde{n}_b = \tilde{n}_b^0 + \tilde{n}_b^1$.

Consider also a *shifted* distribution of publicized contracts $\{n_b^{1,s}(\bar{\gamma})\}_{b=-R}^R$, which is obtained by subtracting a mean price effect $\bar{\gamma}$ to every publicized ($D_t = 0$) contract. That is, $n_b^{1,s}(\bar{\gamma})$ denotes the number of *publicized* contracts with award value p_t such that $(p_t + \bar{\gamma}) \in b$.

Finally, let Δ denote the discrete change in the number of publicized contracts at the discontinuity. Given **A4**, note that this is defined as $\Delta = \delta \cdot \sum_b n_b$.

Propositions

We now make a series of propositions that motivate our estimation method that we label "density analysis" in Section 1.3.

Proposition 3. *There exist some* $(\underline{b}^1, \bar{b}^1)$ *such that* $E[\tilde{n}_b^1] = E[n_b^{s,1}(\bar{\gamma})]$*, for* $\bar{\gamma} = E[\gamma_t]$*,* $b < \underline{b}^1 < 0$ *and* $b > \bar{b}^1 > 0$*. That is, far enough from the threshold, the distribution of* realized *award prices, appropriately shifted to cancel out mean price effects, coincides with the distribution of* ex-ante *award prices for publicized contracts.*

Proposition 4. *There exist some* $(\underline{b}^0, \bar{b}^0)$ *such that* $E[\tilde{n}_b^0] = E[n_b^0]$*, for* $b < \underline{b}^0 < 0$*, and* $b > \bar{b}^0 > 0$*. In other words, far enough from the threshold, the distributions of* ex-ante *and* realized *award prices for non-publicized contracts coincide.*

Corollary 1. $E[\tilde{n}_b] = E[n_b^0 + n_b^{s,1}(\bar{\gamma})]$*, for* $\bar{\gamma} = E[\gamma_t]$*,* $b < \underline{b} = \min\{\underline{b}^0, \underline{b}^1\} < 0$ *and* $b > \bar{b} = \max\{\bar{b}^0, \bar{b}^1\} > 0$.

Proposition 5. $\sum_{b \leq 0}(\tilde{n}_b - n_b) = \sum_{b > 0}(n_b - \tilde{n}_b)$*. This means that the excess mass below the threshold equals the missing mass above the threshold.*

Proposition 6. $\Delta \cdot F_{\gamma'}(x) = E[n_{b_x}^{1,s}(\bar{\gamma}) - \tilde{n}_{b_x}^1]$*, for* $x \in b_x$*,* $b_x \leq 0$*, and* $\gamma' = \gamma - \bar{\gamma}$*.*

Convolution of densities

The key to our propositions stems from characterizing the distribution of observed prices p_t, given the distributions of ex-ante estimates, price effects, and measurement error. Throughout this section, we normalize the price of publicized contracts by subtracting the mean of the price effects. This is for convenience, so that we deal with a mean-zero price effect, but is without loss of generality, as the propositions appropriately adjust for $\bar{\gamma}$ when appropriate.

Consider first the density of publicized contracts, h_p^1. Because observed prices are given by the sum of two independent random variables, ex-ante estimates and price effects (see **A3**), their density is given by the convolution of the densities $f_{\tilde{p}}^1 \equiv f_{\tilde{p}|D=1}$ and f_γ. That is:

$$h_p^1(p_t) = \int_{-\infty}^{\infty} f_{\tilde{p}}^1(p_t - \gamma_t) f_\gamma(\gamma) d\gamma \tag{A.1}$$

On the other hand, using Bayes' rule:

$$f_{\tilde{p}}^1(\tilde{p}_t) = \frac{\tilde{\pi}_D(\tilde{p}_t) \cdot f_{\tilde{p}}(\tilde{p}_t)}{\Pr(D_t = 1)} \tag{A.2}$$

So that (A.1) and (A.2) imply:

$$\begin{aligned} h_p^1(p_t) &= \int_{-\infty}^{\infty} \frac{\tilde{\pi}_D(p_t - \gamma) \cdot f_{\tilde{p}}(p_t - \gamma) \cdot f_\gamma(\gamma)}{\Pr(D_t = 1)} d\gamma \\ &= \int_{-\infty}^{\infty} \frac{(\tilde{\pi}_D^*(p_t - \gamma) + \delta \cdot \mathbf{1}[p_t - \gamma > 0]) \cdot f_{\tilde{p}}(p_t - \gamma) \cdot f_\gamma(\gamma)}{\Pr(D_t = 1)} d\gamma \\ &= \int_{-\infty}^{\infty} \frac{\tilde{\pi}_D^*(p_t - \gamma) \cdot f_{\tilde{p}}(p_t - \gamma) \cdot f_\gamma(\gamma)}{\Pr(D_t = 1)} d\gamma + \int_{-\infty}^{p_t} \frac{\delta \cdot f_{\tilde{p}}(p_t - \gamma) \cdot f_\gamma(\gamma)}{\Pr(D_t = 1)} d\gamma \end{aligned}$$

Or,

$$h_p^1(p_t) \equiv \int_{-\infty}^{\infty} f_{\tilde{p}}^{1*}(p_t - \gamma) \cdot f_\gamma(\gamma) \cdot d\gamma + \int_{-\infty}^{p_t} \Delta(p_t - \gamma) \cdot f_\gamma(\gamma) \cdot d\gamma \tag{A.3}$$

Consider $p_t << 0$, so that $f_\gamma(p_t) \approx 0$. In words, consider a price sufficiently below the threshold, so that the probability that the ex-ante estimate for this contract was above the threshold is negligible. In this case, the second term in Equation (A.3) is zero. On the other hand, $f_{\tilde{p}}^{1*}(p_t - \gamma) = f_{\tilde{p}}^1(p_t - \gamma)$ when $p_t < 0$, so that the first term is the convolution between the densities of $\tilde{p}$ and γ_t. If the former is sufficiently smooth, then adding a mean-zero price effect has no effect on the observed density, and $h_p^1(p_t) = f_{\tilde{p}}^1(p_t)$. It follows that the expected number of contracts with observed price p_t equals the expected number of contracts with ex-ante price estimate equal to p_t. Abandoning the normalization to allow for non-zero average price effects implies that this equality of expectations holds only once observed publicized prices are adjusted by adding the mean of γ. The first part of Proposition 3 follows: for sufficiently low $p_t \in \underline{b}$, $E[\tilde{n}_b^1] = E[n_b^{s,1}(\bar{\gamma})]$, for all $b \leq \underline{b}$.

As we move closer to the threshold from below, the second term in Equation (A.3) becomes positive. This corresponds to the excess mass of contracts, relative to the counterfactual density of the first term. Intuitively, this term is given by the mass of contracts with ex-ante estimate to the right of the threshold that receive a sufficiently high price effect so as to end up at the left of it. This is what allows us to identify F_γ in Proposition 6. Consider $p_t = x$ closely below the threshold, so that $\Delta(x - \gamma) \approx \Delta$. With a constant Δ, it immediately follows that $\Delta \cdot F_\gamma(p_t) = h_p^1(p_t) - f_{\tilde{p}}^1(p_t)$.

A symmetric argument can be given for p_t closely above the threshold. In this case, the second term becomes the missing mass of the observed density $h_p^1(p_t)$, relative to the counterfactual density of $\tilde{p}$. Once we get to a high enough value of $p_t >> 0$, once again $f_\gamma(p_t)$ goes to zero and this missing mass disappears. Observed and counterfactual densities converge, which completes Proposition 3: for sufficiently high $p_t \in \underline{b}$, $E[\tilde{n}_b^1] = E[n_b^{s,1}(\bar{\gamma})]$, for all $b > \overline{b}$.

The argument for non-publicized contracts is directly analogous. Observed awards are the sum of unobserved ex-ante estimates $\tilde{p}$ and a mean-zero error term ξ. This error term only generates a discrepancy between h_p^0 and f_p^0 when the latter is not smooth, which happens only at the threshold. Proposition 4 follows: for $p_t << 0$ and $p_t >> 0$, the two densities coincide.

All this discussion ignored the potential effect of bunching responses. However, strategic bunching does not affect any of the aforementioned results. This is because of **A5**: bunching responses occur only within a window around the threshold. Therefore, all of our arguments remain unchanged, as long as $b_H \leq \overline{b}$, where $p_H \in b_H$.

Finally, Proposition 5 follows directly from the fact that our model assumes no extensive margin responses. Contracting officers can avoid the mandate via bunching responses, but still need to complete the purchase. We think this assumption is natural for this setting, so that the overall number of observed and counterfactual contracts needs to coincide.

Density Analysis: Estimation of Price Effects and Counterfactual Densities

We know explain our density analysis estimation method in detail, building on the Propositions of the previous section.

Step 1

Our method starts from the observation that, relative to ex-ante prices, linear price effects will impact the distribution of publicized contracts in two ways: (i) they will shift the full distribution to the left by $E[\gamma_t]$; and (ii) they will smooth out the discontinuity in the distribution around the threshold, because of $V(\gamma_t)$ (see Figure 1.2 (d)).

Suppose that we knew the true value of mean price effects $E[\gamma_t] \equiv \bar{\gamma}$. From the observed frequency distribution of publicized contracts $\{n_b^1\}$, we can simply undo the first impact of price effects by shifting this distribution back to the right. That is, we construct the *shifted* distribution $\{n_b^{1,s}(\bar{\gamma})\}$, which is obtained by adding the value of $\bar{\gamma}$ to the price award of every publicized contract. If the number of contracts is large, the shifted distribution should coincide with the unobserved distribution of ex-ante prices $\{\tilde{n}_b^d\}$, except near the threshold.

On the other hand, a similar argument can be made for non-publicized contracts, given the assumption that bunching responses are local to the threshold (**A4**). Except for a window around the threshold where bunching responses manifest, the observed distribution $\{n_b^0\}$ should coincide with the unobserved distribution $\{\tilde{n}_b^d\}$ (see Figure 1.2 (c)).

This intuition is supported by Propositions 3 and 4. Once we get "far enough" from the threshold, the distribution of non-publicized awards and the *appropriately shifted* distribution of publicly solicited awards should coincide with the latent distributions of ex-ante prices. In particular, we have that: $n_b^0 + n_b^{1,s}(\bar{\gamma}) \approx \tilde{n}_b^0 + \tilde{n}_b^1 = \tilde{n}_b$ for b sufficiently far from 0. On the contrary, close to the threshold we have $n_b^0 + n_b^{1,s}(\bar{\gamma}) \neq \tilde{n}_b$ due to the effects of bunching and the variance in price effects.

Finally, because we know that the unobserved distribution $\{\tilde{n}_b\}$ should be smooth everywhere due to **A1**, we can use a standard bunching estimation procedure (Chetty, Friedman, and Saez, 2013; Kleven and Waseem, 2013) to infer the shape of it around the threshold. This means fitting a polynomial function through our constructed distribution $\{n_b^0+n_b^{1,s}(\bar{\gamma})\}$, ignoring the contribution of the bins close to the threshold.

More concretely, we estimate the following specification:

$$\left[n_b^0 + n_b^{1,s}(\widehat{\bar{\gamma}})\right] = \sum_{x=0}^{Q} \alpha_x \cdot b^x + \sum_{j=\underline{b}}^{\bar{b}} \gamma_j \cdot \mathbf{1}[b=j] + \nu_b, \qquad \text{for } b = \{-R, ..., R\} \tag{A.4}$$

and obtain fitted values:

$$\widehat{n_b} = \sum_{x=0}^{Q} \widehat{\alpha_x} \cdot b^x \qquad \text{for } b = \{-R, ..., R\}.$$

Now, this discussion started by *assuming* that we knew the value of the mean price effect $\bar{\gamma}$. Yet, in practice, this is the main unknown parameter that we seek to recover. So in order to estimate it, we rely on the integration constraint of Proposition 5: $\sum_{b=-R}^{R}(n_b^0 + n_b^{1,s}(\bar{\gamma})) = \sum_{b=-R}^{R} \widehat{n_b}$. As the intuition from Appendix Figure 1.3 shows, the integration constraints will bind only when we shift the distribution of publicized contracts according to the right value of $\bar{\gamma}$. We, therefore, start from an initial guess of $\widehat{\bar{\gamma}}$, and iterate until we find a value such that the constraint is satisfied.

For the implementation, we choose the following parameters. We use a fifth-degree polynomial, i.e. $Q = 5$. We use bins of constant width of 0.01 log-points. This implies bins of roughly \$250 at the discontinuity. Indeed, bin $b = 0$ includes all contracts with price greater than \$24,751[3] and smaller than or equal to \$25,000. Our estimation is performed on a total set of 150 bins centered around zero, from -0.75 to 0.75. In dollar terms, this corresponds to contracts between \$11,809 and \$52,925. The excluded window for step 1 is symmetric, excluding 12 bins below zero and 12 bins above. In dollar terms, the excluded window consists of contracts between \$22,173 and \$28,187.

[3] $\log(x) - \log(25,000) = 0.01 \iff x = 25,000 \cdot \exp(-0.01)$

Step 2

The second step seeks to estimate separate counterfactual distributions by publicity status, i.e. $\{\widehat{\tilde{n}^0_b}\}$ and $\{\widehat{\tilde{n}^1_b}\}$. To do this, we can go back to the intuition from Figure 1.1, assuming that there are neither price effects nor bunching responses, so that the distributions of ex-ante prices and observed realized prices coincide. In this case, the distributions for treated and control units should be continuous, except at the threshold, where we should see a discontinuous jump in publicized contracts mirrored by a discontinuous dip in non-publicized contracts. Suppose that we knew the size of this change, which we denote as Δ. Knowledge of Δ would allow us to undo these discontinuities by shifting the right part of each distribution vertically. Indeed, the distributions $\{n^0_b + \Delta \cdot \mathbf{1}[b > 0]\}$ and $\{n^1_b - \Delta \cdot \mathbf{1}[b > 0]\}$ should be continuous.

In the presence of bunching and price effects, these vertical shifts will not make the observed distributions continuous. However, just as in the discussion above, price effects and bunching should only affect the distributions within some window around the threshold. So we use this logic again and use a polynomial interpolation to estimate the counterfactual distributions around the threshold.

First, we construct distributions that are vertically shifted above the threshold: $\{n^0_b + \Delta \cdot \mathbf{1}[b > 0]\}^R_{b=-R}$ and $\{n^{1,s}_b(\widehat{\gamma}) - \Delta \cdot \mathbf{1}[b > 0]\}^R_{b=-R}$. We then apply the same interpolation method as before for each of the two distributions. That is, we separately estimate the following two specifications:

$$\left(n^0_b + \Delta \cdot \mathbf{1}[b > 0]\right) = \sum_{x=0}^{Q} \alpha^0_x \cdot b^x + \sum_{j=\underline{b^0}}^{\overline{b^0}} \gamma^0_j \cdot \mathbf{1}[b = j] + \nu^0_b, \quad \text{for } b = \{-R, ..., R\} \tag{A.5}$$

$$\left(n^{1,s}_b(\widehat{\gamma}) - \Delta \cdot \mathbf{1}[b > 0]\right) = \sum_{x=0}^{Q} \alpha^1_x \cdot b^x + \sum_{j=\underline{b^1}}^{\overline{b^1}} \gamma^1_j \cdot \mathbf{1}[b = j] + \nu^1_b, \quad \text{for } b = \{-R, ..., R\} \tag{A.6}$$

and compute fitted values ignoring the contribution of the bins within the excluded window:

$$\widehat{n^{*0}_b} = \sum_{x=0}^{Q} \widehat{\alpha^0_x} \cdot b^x, \quad \text{for } b = \{-R, ..., R\}$$

$$\widehat{n^{*1}_b} = \sum_{x=0}^{Q} \widehat{\alpha^1_x} \cdot b^x, \quad \text{for } b = \{-R, ..., R\}$$

Finally, our estimates of the counterfactual distributions do incorporate the discontinuous effect of the policy. We estimate these by re-adding the shift that we originally removed:

$$\widehat{n^0_b} = \widehat{n^{*0}_b} - \Delta \cdot \mathbf{1}[b > 0] \quad \text{for } b = \{-R, ..., R\}$$

$$\widehat{n^1_b} = \widehat{n^{*1}_b} + \Delta \cdot \mathbf{1}[b > 0] \quad \text{for } b = \{-R, ..., R\}$$

Again, this exposition assumes that we know the value of Δ. Since, in practice, this is not directly observed, our method iterates over guesses of $\widehat{\Delta}$. The convergence criterion in this case is based on the fit of the interpolations outside the excluded window. Indeed, if the vertical shift we guess is too low or too high, the polynomial interpolation will fit poorly just outside of the excluded area. Appendix Figure 1.4 shows this intuition graphically.

So, given a guess of $\widehat{\Delta}$, we compute the residuals for each of the two regressions (A.5) and (A.6). We then search over $\widehat{\Delta}$ to minimize:

$$W\left(\widehat{\Delta}\right) = 0.5 \cdot \sum_{b \neq Z^0} \widehat{\nu_b^0}\left(\widehat{\Delta}\right)^2 + 0.5 \cdot \sum_{b \neq Z^1} \widehat{\nu_b^1}\left(\widehat{\Delta}\right)^2 \quad ,$$

where $Z^0 = \{\underline{b^0}, ..., \overline{b^0}\}$ and $Z^1 = \{\underline{b^1}, ..., \overline{b^1}\}$ correspond to the excluded regions.

For step two, we keep the polynomial degree fixed, binning and range fixed as in step 1. However, we change the excluded region for the specification using non-publicized contracts (A.5). The justification of this is that we expect bunching to be concentrated closely below the threshold. Concretely, we choose 5 bins below the threshold and 12 bins above for Z^0 and keep the symmetric window of 12 bins above and below for Z^1.

Step 3

In step 3 we rely on the formula from Proposition 6 and use our estimates from above to compute:

$$\widehat{F_{\gamma'}}(x) = \frac{n_{b_x}^{1,s}(\widehat{\bar{\gamma}}) - \widehat{\tilde{n}_{b_x}^1}}{\widehat{\Delta}}$$

for $x \in b_x$, $b_x \in \{\underline{b}^1, ..., 0\}$, and $\gamma' = \gamma - \widehat{\bar{\gamma}}$. This is straightforward given implementation of steps 1 and 2. We obtain the $F_{\gamma'}$ evaluated at each bin on the lower half of the excluded region Z^1. For values $x < \underline{b}^1$, we impose $F_{\gamma'} = 0$, since below the excluded region there is no longer any influence of price effects. Finally, we then obtain estimates for the rest of the CDF by imposing symmetry, so that $F_{\gamma'}(x) = 1 - F_{\gamma'}(-x)$.

For all of our estimates, we compute standard errors via bootstrap. We sample with replacement from the original distribution of contracts, and implement steps 1 through 3, obtaining a set of estimates $\hat{\theta}$. We repeat this process H times. The standard errors correspond to the empirical standard deviation of $\hat{\theta}^{(h)}$, for $h = \{1, 2, ..., H\}$.

RDD Correction for Price Effects and Measurement Error

Consider again the model described in Section A.4. Observed prices as a function of ex-ante prices are given by:

$$p_t = \tilde{p}_t + (1 - D_t) \cdot \xi_t + D_t \cdot \gamma_t \tag{A.7}$$

where p_t are observed normalized (i.e. logged and re-centered around 0) award prices, $\tilde{p}_t$ are normalized ex-ante prices, $D_t \in \{0, 1\}$ are publicity decisions, γ_t is the *price effect* of publicity, and ξ_t is *measurement error*. Let $\gamma_t \sim F_\gamma(\cdot)$, with $E[\gamma_t] = \mu_\gamma$ and $V[\gamma_t] = \sigma^2_\gamma$. Let $\xi_t \sim F_\xi(\cdot)$, with $E[\xi_t] = 0$ and $V[\xi_t] = \sigma^2_\xi$. Assume $\gamma_t \perp \xi_t \perp \tilde{p}_t$.

To assess the causal impact of D_t on outcomes of interest y_t, we assume a piece-wise linear relationship between expected outcomes and latent ex-ante prices. In particular:

$$E[y_t|\tilde{p}_t] = \mathbf{1}(\tilde{p}_t \leq 0) \cdot (\alpha_0 + \beta_0 \cdot \tilde{p}_t) + \mathbf{1}(\tilde{p}_t > 0) \cdot (\alpha_1 + \beta_1 \cdot \tilde{p}_t) \tag{A.8}$$

For simplicity, we focus on this reduced form relationship, but it would be straightforward to extend it to a two-equation model with a structural equation relating y_t and D_t, and a first-stage equation relating D_t and $\tilde{p}_t$. Our parameters of interest are $(\boldsymbol{\alpha}, \boldsymbol{\beta}) = (\alpha_0, \alpha_1, \beta_0, \beta_1)$. In particular, we focus on $(\alpha_1 - \alpha_0)$, the reduced form effect at the discontinuity.

The problem we face is that we do not observe a sample analog of $E[y_t|\tilde{p}_t]$, but rather of $E[y_t|p_t]$. Our "naive RDD" coefficients correspond to an estimate of $(\lim_{p\to 0^+} E[y_t|p] - \lim_{p\to 0^-} E[y_t|p])$, which in general will not be equal to $(\alpha_1 - \alpha_0) = (\lim_{\tilde{p}\to 0^+} E[y_t|\tilde{p}] - \lim_{\tilde{p}\to 0^-} E[y_t|\tilde{p}])$. Here we propose an alternative estimator of $(\alpha_1 - \alpha_0)$ based on the following proposition.

Proposition 7. *Expected outcomes conditional on observed award prices $E[y_t|p_t]$ can be expressed as an explicit linear function of the structural parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta})$, as well as other variables that we can directly observe or estimate. In particular:*

$$E[y_t|p_t] = \alpha_0 \cdot \psi_1(p_t) + \beta_0 \cdot \psi_2(p_t) + \alpha_1 \cdot \psi_3(p_t) + \beta_1 \cdot \psi_4(p_t) \quad ,$$

where $\psi_k(\cdot)$, $k \in \{1,2,3,4\}$ are explicit functions of observed prices (p_t), observed treatment probabilities at a given price $(\pi_D(p_t))$, and moments of the distributions of price effects and measurement error evaluated at a given price $(F_\gamma(p_t), F_\xi(p_t))$.

Below we derive the explicit expressions for each ψ_k. We then compute these using our data and the estimate $\widehat{F_\gamma}(p_t)$ that we obtained from the density analysis. We also assume no measurement error, so that $\xi_t = 0$ for all t. However, the formulas we derive are general, allowing for any arbitrary distribution of measurement error. Once we compute these estimates $\widehat{\psi_k}(p_t)$, we use the equation in Proposition 7 to estimate $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ by OLS. We are particularly interested in $(\hat{\alpha}_1^{OLS} - \hat{\alpha}_0^{OLS})$, which we then directly compare to the "naive RDD" reduced form coefficients.

Proof of Proposition 7

We now derive the explicit expression for $E[y_t|p_t]$. First, we use the Law of Total Probability to write:

$$E[y_t|p_t] = \underbrace{E[y_t|p_t, \tilde{p}_t \leq 0]}_{\Lambda_1} \cdot \underbrace{\Pr(\tilde{p}_t \leq 0|p_t)}_{\Lambda_2} + \underbrace{E[y_t|p_t, \tilde{p}_t > 0]}_{\Lambda_3} \cdot \underbrace{\Pr(\tilde{p}_t > 0|p_t)}_{\Lambda_4} \tag{A.9}$$

For each Λ_k, $k \in \{1,2,3,4\}$, we find an expression that depends only on magnitudes that we can directly observe or estimate.

We start with Λ_2:

$$\begin{aligned}
\Lambda_2 &= \Pr(\tilde{p}_t \leq 0|p_t) \\
&= \Pr(\tilde{p}_t \leq 0|p_t, D_t = 0) \cdot \Pr(D_t = 0|p_t) + \Pr(\tilde{p}_t \leq 0|p_t, D_t = 1) \cdot \Pr(D_t = 1|p_t) \\
&= \Pr(p_t - \xi_t \leq 0|p_t) \cdot [1 - \pi_D(p_t)] + \Pr(p_t - \gamma_t \leq 0|p_t, D_t = 1) \cdot \pi_D(p_t) \\
&= [1 - F_\xi(p_t)] \cdot [1 - \pi_D(p_t)] + [1 - F_\gamma(p_t)] \cdot \pi_D(p_t) \\
&\equiv \Lambda_2(p_t, \pi_D(p_t), F_\gamma(p_t), F_\xi(p_t), \boldsymbol{\alpha}, \boldsymbol{\beta})
\end{aligned} \tag{A.10}$$

Similarly for Λ_4:

$$\begin{aligned}
\Lambda_4 &= \Pr(\tilde{p}_t \geq 0|p_t) \\
&= \Pr(\tilde{p}_t \geq 0|p_t, D_t = 0) \cdot \Pr(D_t = 0|p_t) + \Pr(\tilde{p}_t \geq 0|p_t, D_t = 1) \cdot \Pr(D_t = 1|p_t) \\
&= \Pr(p_t - \xi_t \geq 0|p_t) \cdot [1 - \pi_D(p_t)] + \Pr(p_t - \gamma_t \geq 0|p_t, D_t = 1) \cdot \pi_D(p_t) \\
&= F_\xi(p_t) \cdot [1 - \pi_D(p_t)] + F_\gamma(p_t) \cdot \pi_D(p_t) \\
&\equiv \Lambda_4(p_t, \pi_D(p_t), F_\gamma(p_t), F_\xi(p_t), \boldsymbol{\alpha}, \boldsymbol{\beta})
\end{aligned} \tag{A.11}$$

For Λ_1 and Λ_3, the analysis is slightly more complicated. First, observe that:

$$\begin{aligned}
\Lambda_1 &= E[y_t|p_t, \tilde{p}_t \leq 0]) \\
&= E[\alpha_0 + \beta_0 \cdot \tilde{p}_t|p_t, \tilde{p}_t \leq 0] \\
&= \alpha_0 + \beta_0 \cdot E[\tilde{p}_t|p_t, \tilde{p}_t \leq 0] \\
&= \alpha_0 + \beta_0 \cdot \{E[\tilde{p}_t|p_t, \tilde{p}_t \leq 0, D_t = 1] \cdot \Pr(D_t = 1|p_t, \tilde{p}_t \leq 0) \\
&\qquad\qquad + E[\tilde{p}_t|p_t, \tilde{p}_t \leq 0, D_t = 0] \cdot \Pr(D_t = 0|p_t, \tilde{p}_t \leq 0)\} \\
&= \alpha_0 + \beta_0 \cdot \{E[p_t - \gamma_t|p_t, \tilde{p}_t \leq 0, D_t = 1] \cdot \Pr(D_t = 1|p_t, \tilde{p}_t \leq 0) \\
&\qquad\qquad + E[p_t - \xi_t|p_t, \tilde{p}_t \leq 0, D_t = 0] \cdot \Pr(D_t = 0|p_t, \tilde{p}_t \leq 0)\} \\
&= \alpha_0 + \beta_0 \cdot \{(p_t - E[\gamma_t|\gamma_t \geq p_t, p_t]) \cdot \Pr(D_t = 1|p_t, \tilde{p}_t \leq 0) \\
&\qquad\qquad + (p_t - E[\xi_t|\xi_t \geq p_t, p_t]) \cdot \Pr(D_t = 0|p_t, \tilde{p}_t \leq 0)\}
\end{aligned}$$

$\Longleftrightarrow$

$$\begin{aligned}
\Lambda_1 = \alpha_0 + \beta_0 \cdot p_t + \beta_0 \cdot \{&E[\gamma_t|\gamma_t \geq p_t, p_t]) \cdot \Pr(D_t = 1|p_t, \tilde{p}_t \leq 0) \\
&- E[\xi_t|\xi_t \geq p_t, p_t]) \cdot \Pr(D_t = 0|p_t, \tilde{p}_t \leq 0)\}
\end{aligned} \tag{A.12}$$

Now, applying Bayes' rule to $\Pr(D_t = 0|p_t, \tilde{p}_t \leq 0)$:

$$\begin{aligned}
\Pr(D_t = 0|p_t, \tilde{p}_t \leq 0) &= \frac{\Pr(\tilde{p}_t \leq 0|D_t = 0, p_t) \cdot \Pr(D_t = 0|p_t)}{\Pr(\tilde{p}_t \leq 0|p_t)} \\
&= \frac{\Pr(\tilde{p}_t \leq 0|D_t = 0, p_t) \cdot \Pr(D_t = 0|p_t)}{\Lambda_2} \\
&= \frac{\Pr(p_t - \xi \leq 0|p_t) \cdot [1 - \pi_D(p_t)]}{\Lambda_2} \\
&= \frac{[1 - F_\xi(p_t)] \cdot [1 - \pi_D(p_t)]}{\Lambda_2}
\end{aligned} \tag{A.13}$$

And, therefore,

$$\begin{aligned}
\Pr(D_t = 1|p_t, \tilde{p}_t \leq 0) &= 1 - \Pr(D_t = 0|p_t, \tilde{p}_t \leq 0) \\
&= \frac{[1 - F_\gamma(p_t)] \cdot \pi_D(p_t)}{\Lambda_2}
\end{aligned} \tag{A.14}$$

Combining (A.12), (A.13) and (A.14) implies:

$$\begin{aligned}
\Lambda_1 &= \alpha_0 + \beta_0 \left[p_t + \frac{E[\gamma_t|\gamma_t \geq p_t, p_t] \cdot [1 - F_\gamma(p_t)] \cdot \pi_D(p_t) - E[\xi_t|\xi_t \geq p_t, p_t] \cdot [1 - F_\xi(p_t)] \cdot [1 - \pi_D(p_t)]}{\Lambda_2} \right] \\
&\equiv \Lambda_1(p_t, \pi_D(p_t), F_\gamma(p_t), F_\xi(p_t), \boldsymbol{\alpha}, \boldsymbol{\beta})
\end{aligned} \tag{A.15}$$

Analogous calculations yield the following expression for Λ_3:

$$\begin{aligned}
\Lambda_3 &= \alpha_1 + \beta_1 \left[p_t + \frac{E[\gamma_t|\gamma_t \leq p_t, p_t] \cdot F_\gamma(p_t) \cdot \pi_D(p_t) - E\left[\xi_t|\xi_t \leq p_t, p_t\right] \cdot F_\xi(p_t) \cdot [1 - \pi_D(p_t)]}{\Lambda_4} \right] \\
&\equiv \Lambda_3(p_t, \pi_D(p_t), F_\gamma(p_t), F_\xi(p_t), \boldsymbol{\alpha}, \boldsymbol{\beta})
\end{aligned} \tag{A.16}$$

Finally, combining (A.9), (A.10), (A.11), (A.15), and (A.16), we obtain:

$$E[y_t|p_t] = \alpha_0 \cdot \psi_1(p_t) + \beta_0 \cdot \psi_2(p_t) + \alpha_1 \cdot \psi_3(p_t) + \beta_1 \cdot \psi_4(p_t)$$

where:

$$\begin{aligned}
\psi_1(p_t) &= [1 - F_\xi(p_t)] \cdot [1 - \pi_D(p_t)] + [1 - F_\gamma(p_t)] \cdot \pi_D(p_t) \\
\psi_2(p_t) &= \psi_1(p_t) \cdot p_t + E[\gamma_t|\gamma_t \geq p_t, p_t] \cdot [1 - F_\gamma(p_t)] \cdot \pi_D(p_t) - E[\xi_t|\xi_t \geq p_t, p_t] \cdot [1 - F_\xi(p_t)] \cdot [1 - \pi_D(p_t)] \\
\psi_3(p_t) &= F_\xi(p_t) \cdot [1 - \pi_D(p_t)] + F_\gamma(p_t) \cdot \pi_D(p_t) \\
\psi_4(p_t) &= \psi_3(p_t) \cdot p_t + E[\gamma_t|\gamma_t \leq p_t, p_t] \cdot F_\gamma(p_t) \cdot \pi_D(p_t) - E\left[\xi_t|\xi_t \leq p_t, p_t\right] \cdot F_\xi(p_t) \cdot [1 - \pi_D(p_t)]
\end{aligned}$$

Accounting for Bunching

A standard test for the validity of the RDD framework consists on verifying the continuity of the density of the running variable around the threshold. If the running variable is not distributed smoothly around the cutoff, then it is said to be "manipulated". In recent work, Gerard, Rokkanen, and Rothe (2020) show that, while point identification of causal effects is infeasible in this case, it is possible to obtain sharp bounds on the effects of interest.

In their model, the extent of manipulation can be quantified as the excess bunching in the density of the running variable below the threshold. While one cannot identify which are the units below the threshold that are manipulating, the excess bunching π_B tells us what share of the observed units are in this group. Bounds on treatment effects are then computed by excluding a share π_B of the observations below the threshold, in ways that yield the most extreme values for the estimate.

This process can be quite involved in general, since one does not know the treatment assignment of the units that manipulate. This transforms the computation of the bounds in an optimization problem, searching for the worst- and best-case scenarios in terms of how outcomes are distributed across treatment groups below the threshold.

However, our setting allows us to make a behavioral assumption that tremendously simplifies the problem. In particular, our model assumes that all units that manipulate the ex-ante price to bunch below the threshold, successfully avoid the publicity mandate.[4] Therefore, our model implies that the share π_B of units that manipulate all belong to the control group ($D_t = 0$). Bounds on treatment effects are straightforwardly obtained in this case, by simply chopping the tails of the distribution of outcomes Y_t below the threshold for units in the control group.

In practice, we implement this procedure as follows. For each bin b closely below the threshold:

1. Compute the excess bunching in the control group, as $BUNCH_b = (n_b^0 - \widehat{\tilde{n}_b^0})$, obtained from our density analysis.
2. Sort control units according to the outcome variable Y_b^0.
3. Drop the $BUNCH_b$ units with the highest value of Y_b^0. Compute treatment effects. This yields the lower bound.
4. Drop the $BUNCH_b$ units with the lowest value of Y_b^0. Compute treatment effects. This yields the upper bound.

[4] This corresponds to a special case of their more general model. The authors explicitly discuss this special case in their Appendix C.3.

Appendix B

Appendix: Competition under Incomplete Contracts and the Design of Procurement Policies II: Structural Estimates

B.1 Additional Figures

Figure B.1.1: Complexity Distribution

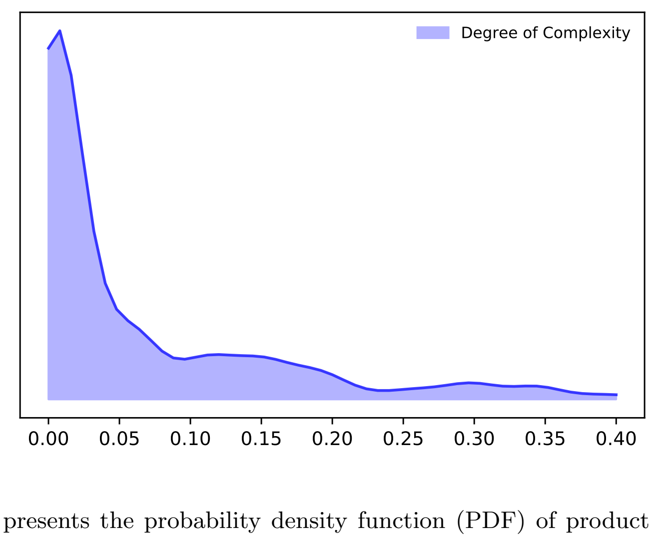

Notes: This figure presents the probability density function (PDF) of product complexity. Even though there's wide heterogeneity on the degree of complexity, the bulk of contracts in our sample have relatively low levels of complexity. The degree of complexity is defined as the log of the product's average overruns, and it is calculated on all contracts for the same product category that are smaller than the regulation threshold ($25,000). The plotted distribution of log costs is smoothed using a kernel.

Figure B.1.2: Model Fit

(a) Publicity Status

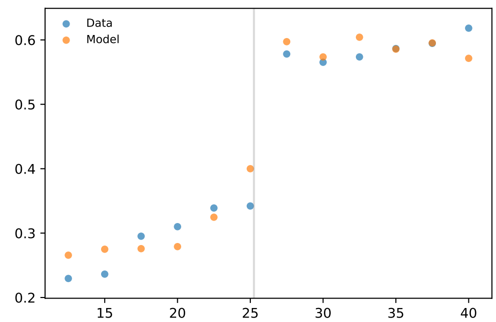

(b) Price Density

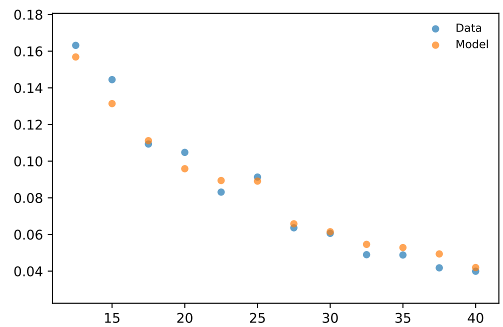

(c) Number of Bidders

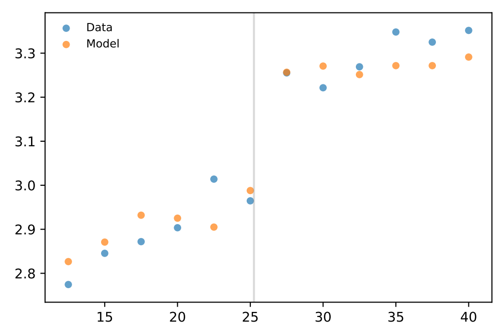

(d) Probability of a Local Winning

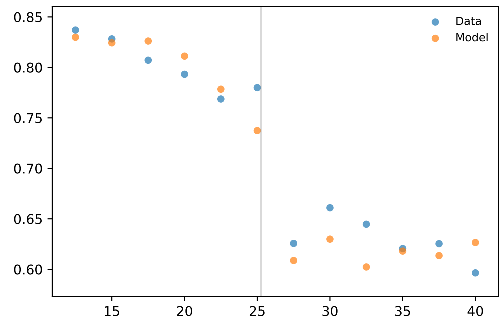

(e) Overruns

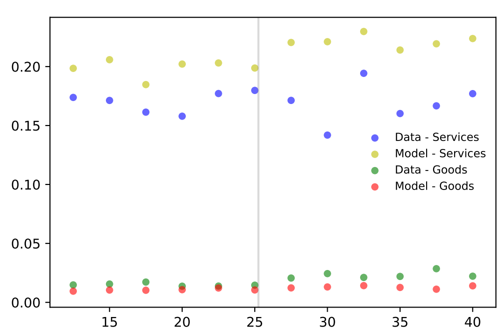

(f) Any Overruns

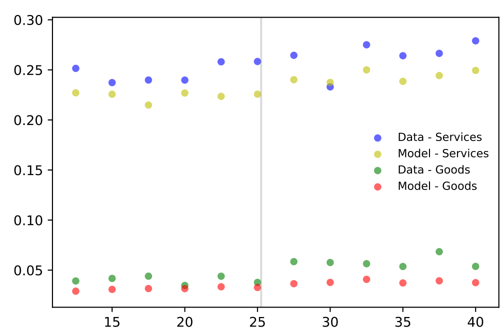

Notes: This figure presents the model fit, based on a simulated method of moments estimation. In each panel, relevant outcome variables relative to the awarding price. Actual data are presented in blue, while model-based simulated data is presented in orange. Panel (a) presents the density of contract prices, Panel (b) the fraction of publicized contracts, Panel (c) the number of actual bidders, Panel (d) fraction awarded to local contractors, Panel (e) and (f) show the average overrun and the chances of having any overrun. The last two panels separate goods from services. The simulated outcomes simulate unobservables building upon actual data. The model is estimated using 24,135 observations, the simulated methods expand each observation multiple times.

Figure B.1.3: Participation

(a) Composition of Bidders

(b) Identity of the Winner

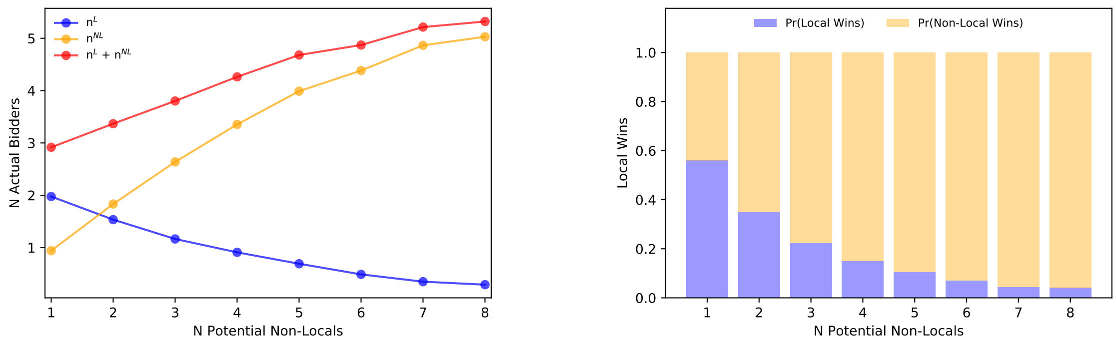

Notes: This figure presents the participation decisions and the subsequent winner identity as a function of the number of potential bidders. Panel (a) the number of actual bidders of each group, the panel (b) displays the average probability of awarding local bidders. The higher the number of potential non-locals, the less likely that locals participate and win. These features connect directly with the fact that locals have substantially higher participation costs; thus, in equilibrium, reductions on predicted utility due to increased competition discourage their participation. Both figures were generated keeping constant (at the mean) the number of potential locals.

B.2 Additional Tables

Table B.2.1: Variable Description Model

	Model Sample	Full Sample	Diff
Variables:			
Publicized in FBO	0.373	0.274	0.099
Award Amount	21.178	20.627	0.551
Number of Offers	3.002	3.098	-0.096
Overruns (relative)	1.117	1.088	0.029
Service	0.375	0.308	0.067
Mean Overruns Prod Cat	0.089	0.071	0.018
Awarded in September	0.249	0.262	-0.013
log Duration	3.976	3.811	0.165
Bidders' Classification			
Local is Awarded	0.754	-	-
N Potential Local Bidders	6.078	-	-
N Potential Non-Local Bidders	3.339	-	-
Number of Observations	24,135	103,899	

Notes: This first column describes the mean of the variables included in the model estimation. The second column shows variables mean but over the full sample. The third column shows the differences between these two means. The model's sample corresponds to the subset of contracts over which we could identify the number of potential local and non-local bidders. We restrict the analysis to buyer-product combinations that meet two conditions: at least four contracts were awarded between 2013 and 2019, and not all nor none were publicized.

B.3 Model Estimation Details

Estimation

Denote the target moments by m_n as a vector of moments from the data. The simulated moments are denoted by $m_s(\theta)$. The depends on the parameters $\theta \in \Theta \subset \mathbb{R}^P$. The estimator minimizes the standard distance metric:

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

Where W_n is the weighting matrix, which is chosen using the standard two-step approach. Letting $M_s(\theta)$ be the $(P \times J)$ Jacobian matrix of the vector of simulated moments; under standard regularity assumptions, we have:

$$\sqrt{n}\left(\hat{\theta} - \theta_0\right) \xrightarrow{d} N\left(0, \left(1 + \frac{1}{s}\right)(M'WM)^{-1}M'W\Omega W'M(M'WM)^{-1}\right) \tag{B.1}$$

where W is the probability limit of W_n, M is the probability limit of $M_s(\theta_0)$, and Ω is the asymptotic variance of m_n (Pakes and Pollard, 1989). The vector of parameters is: $\theta = (\alpha^k, \nu^k, \tau^k, \gamma^k, \xi^k, \vec{\beta}, \zeta, \vec{\sigma})$.

Standard Errors

We compute standard errors using the asymptotic variance formula given by (B.1). The variance-covariance matrix of $\hat{\theta}$ is:

$$V(\hat{\theta}) = \frac{1}{n}\left(1 + \frac{1}{s}\right)(\hat{M}'W\hat{M})^{-1}\hat{M}'W\hat{\Omega}W'\hat{M}(\hat{M}'W\hat{M})^{-1}$$

Where $\hat{\Omega}$ is estimated via bootstrap: re-sampling contracts with replacement from the original data, and recompute the smoothed vector of moments, repeating this process 500 times. $\hat{\Omega}$ is the sample variance of these 500 vectors. $\hat{M}$ is the numeric derivative of the SMM objective function (2.3) evaluated at $\hat{\theta}$.

Minimization

We keep constant the underlying random draws throughout the minimization of the objective function. Nonetheless, the simulated objective is not continuous with respect to θ. Thus, We leverage the stochastic optimization algorithm *Differential Evolution* (Storn and Price, 1997) to perform the objective minimization. This algorithm does not rely on gradient methods, and given its heuristic approach for minimizing possibly nonlinear and non-differentiable continuous space functions, it is robust to poorly behaved objectives.

Moments

We use three sets of target moments.

- First set of moments,
  - $\vec{m}_{11} = \mathbb{E}[x_t^{y}{}'(y_t)]$ and $\vec{m}_{12} = \mathbb{E}[x_t^{y}{}'(y_t)^2]$, where y_t = *log winning bid, number of bidders, wins local, log overruns,* and *contract is publicized*, and $x_t^y = (1, \vec{x}_t^y)$ = covariates associated with the generation of outcome y
- Second set:
  - $\vec{m}_2 = \mathbb{E}[y_t | B_t \in (B^l, B^{l+1})]$, for $l \in \{1, \ldots, L-1\}$, where y_t = *number of bidders, wins local, log overruns,* and *contract is publicized.* We separate these moments based on goods and services, and partition the domain of contract prices in bins of width \$1,000
- Third set of moments:
  - $\vec{m}_3 = \mathbb{E}[\mathbb{1}\{b_t \in (b^l, b^{l+1})\}]$, for $l \in \{1, \ldots, L-1\}$. This set of moments correspond to the normalized frequencies on the relevant window of contract prices. The bin width is \$1,000.

As a result we use 357 moments to estimate 37 parameters.

Appendix C

Appendix: Slippery Fish: Enforcing Regulation when Agents Learn and Adapt

C.1 Appendix Figures on the Research Context

Fishermen Villages

Figure C.1.1: Fishermen Village (Caleta)

Outdoor Markets

Figure C.1.2: Examples of Ferias

Figure C.1.3: Example of a Circuit

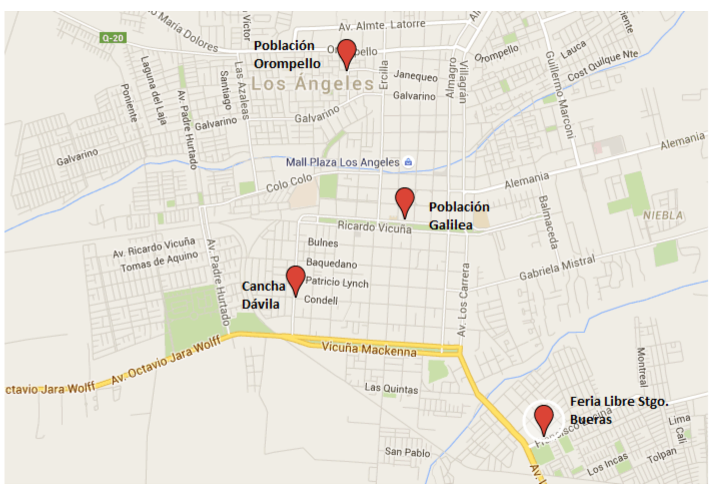

Figure C.1.3 maps the four ferias that compound one circuit of the city of Los Angeles, VII region.

Figure C.1.4: Map of Circuits and Caletas

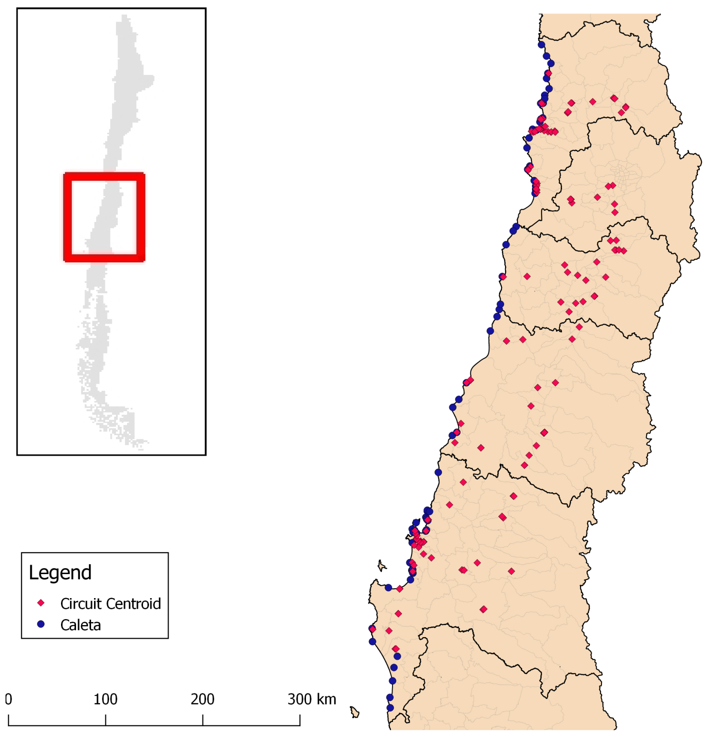

Interventions

Figure C.1.5: Flyers

The figure C.1.5 shows the two types of flyers distributed during the ban period. The message of the one in the right says, "In September respect the Ban", the one in the right says "This month respect the Ban".

Figure C.1.6: Letter to Consumers

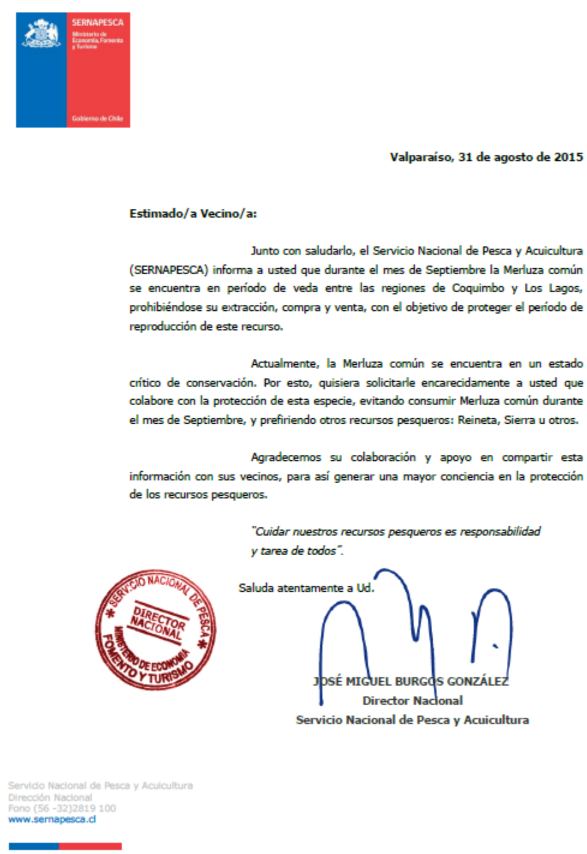

Valparaíso, 31 de agosto de 2015

Estimado/a Vecino/a:

Junto con saludarlo, el Servicio Nacional de Pesca y Acuicultura (SERNAPESCA) informa a usted que durante el mes de Septiembre la Merluza común se encuentra en período de veda entre las regiones de Coquimbo y Los Lagos, prohibiéndose su extracción, compra y venta, con el objetivo de proteger el período de reproducción de este recurso.

Actualmente, la Merluza común se encuentra en un estado crítico de conservación. Por esto, quisiera solicitarle encarecidamente a usted que colabore con la protección de esta especie, evitando consumir Merluza común durante el mes de Septiembre, y prefiriendo otros recursos pesqueros: Reineta, Sierra u otros.

Agradecemos su colaboración y apoyo en compartir esta información con sus vecinos, para así generar una mayor conciencia en la protección de los recursos pesqueros.

"Cuidar nuestros recursos pesqueros es responsabilidad y tarea de todos".

Saluda atentamente a Ud.

JOSÉ MIGUEL BURGOS GONZÁLEZ
Director Nacional
Servicio Nacional de Pesca y Acuicultura

Servicio Nacional de Pesca y Acuicultura
Dirección Nacional
Fono (56 -32)2819 100
www.sernapesca.cl

The figure C.1.6 shows the letter distributed to households during September 2015. The letter, signed by Sernapesca's director, informs about the September ban and the fact that hake's conservation is threatened because of overfishing.

Figure C.1.7: Examples of Neighborhood Treatment Assignment

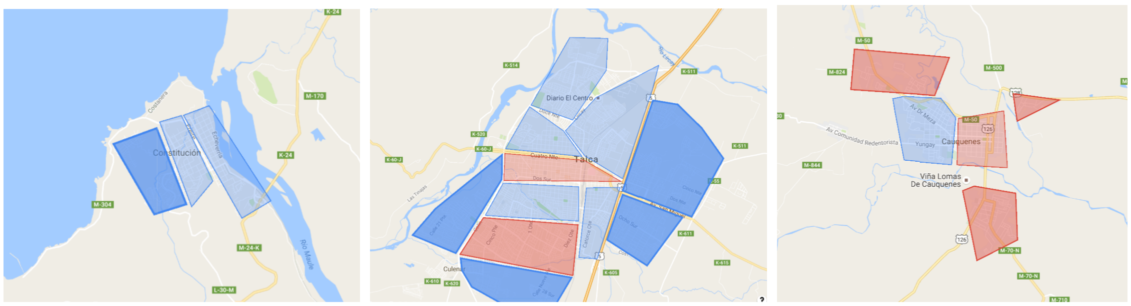

The 48 most populated comunas were divided randomly into three levels of saturation: high, low and zero. Based on the level of saturation, the information campaign was assigned at the neighborhood level. The figure C.1.7 shows the map of three different comunas: The comuna in the left didn't receive information campaign, the one in the center received low level of saturation, the one in the right received high level of saturation. In red, those neighborhoods assigned to receive the information campaign.

Figure C.1.8 describes our interventions and data collection activities along the supply chain for fish. We collected data through surveys and mystery shoppers visits on seven occasions between August 2015 and August 2016.

Figure C.1.8: Interventions and Data Collection at different Points along the Fish Supply Chain

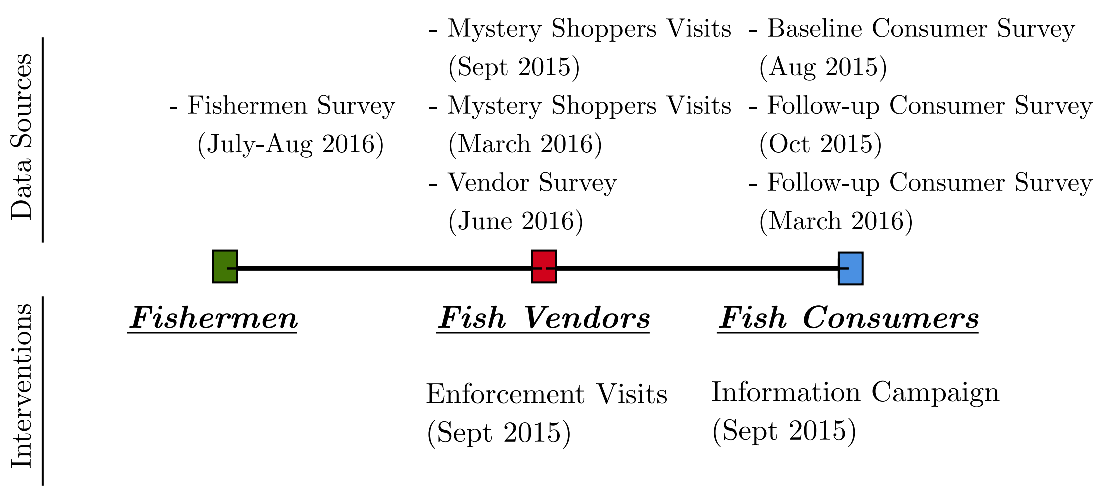

C.2 Appendix Figures on Results

Adoption of Defensive Actions

Figure C.2.1 describes the unconditional probability of selling hake defensively (either frozen or hidden) by week. The probability of selling defensively increases over time, while the overall probability of selling hake decreases substantially along the month. The conditional probability is presented in C.2.1.

Figure C.2.1: Adoption of Defensive Actions

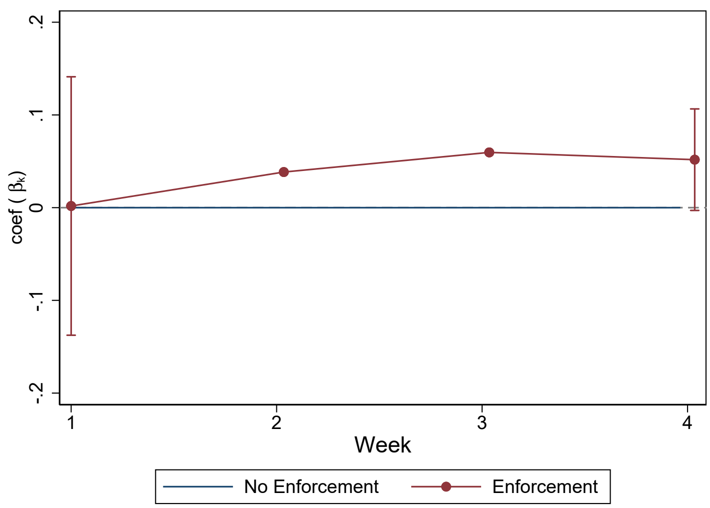

Figure C.2.1 describes the unconditional probability of selling hake either frozen or hidden. The coefficients were obtained from an OLS regression in which the treatment assignment interacts with weekly dummies. We include strata fixed effects and cluster at the circuit level. The "No Enforcement" category is the omitted category and includes observations assigned to the control group and the information campaign. To facilitate the interpretation, we only present the confidence intervals associated with weeks one and four.

Number of Stalls

Figure C.2.2 describes the average number of fish stalls in the first and second half of the month. It shows that the average number of fish stalls does decrease in the markets randomly assigned to the enforcement treatment, especially during the second half of September.

Figure C.2.2: Number of stalls in Feria by Treatment Assignment

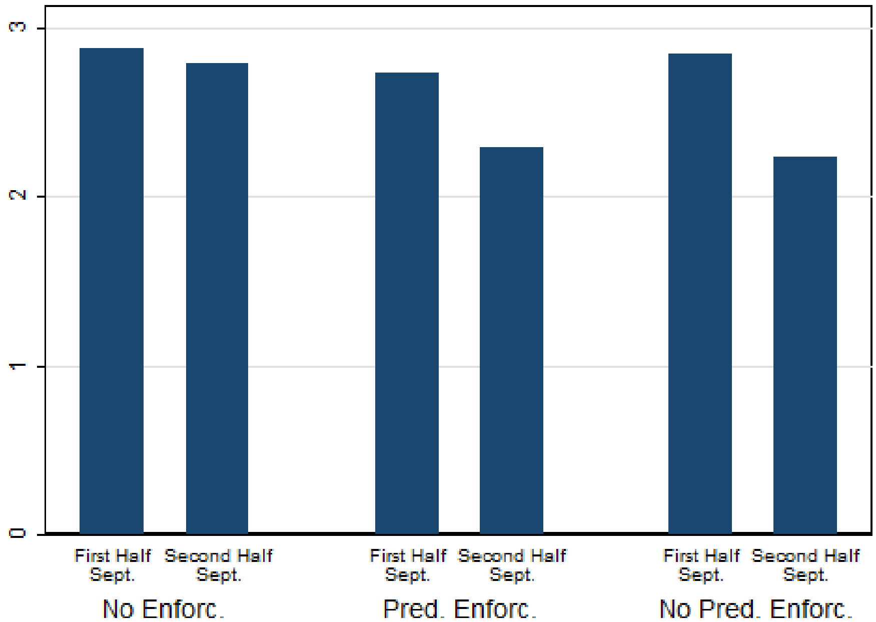

This figure shows the average number of stalls in each feria, separately for the first and the second half of the month. Markets assigned to receive enforcement showed a decrease in the number of stalls between the first and the second half of the month.

Prices

Figure C.2.3 describes the week-by-week evolution of (log) prices for hake and pomfret during the ban. Both rates are normalized to their levels in the first week. It shows that the price of hake increased week-to-week, throughout the ban period; the hake price in the fourth week is 40% higher than the first price. Pomfret prices fell by 10% in the second week, and that lower prices remained stable after that.

Figure C.2.3: Log Prices of Fish During the Ban

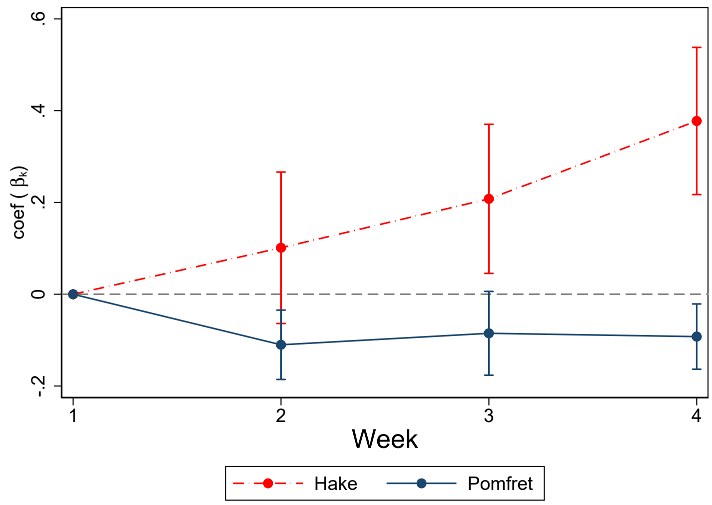

This figure shows the evolution of log prices for hake and pomfret, using the first week as a reference. The price of hake continuously increased over the course of September 2015. Hake was 40% more expensive by the fourth week relative to the first week. The price of pomfret decreased around 10% after the first week.

C.3 Appendix Tables

Descriptive Statistics

Data Collected by Mystery Shoppers

During September 2015, the mystery shoppers interacted with fish-vendors 908 times. The table C.3.1 describes observable characteristics of the stalls visited and the vendors. In general, each stall was operated by one person. Mostly man, and based on mystery shoppers' guess, 47-year-old. The type of weight used informs about the formality of the stall; digital weights are more precise and expensive.

Table C.3.1: Fish Vendors in Ferias

Variable	Mean	SD	Min	Max	N
Number of Vendors per Stall	1.089	0.326	1	4	883
Proportion Female Fish Vendor	0.425	0.479	0	1	882
Age Vendor	47.438	10.126	19	75	876
Prices Visibly Listed	0.242	0.429	0	1	908
Type of Weight					
No Weight	0.262	0.440	0	1	848
Mechanical Weight	0.410	0.492	0	1	848
Digital Weight	0.308	0.462	0	1	848

Notes: This table presents observable characteristics of fish stalls visited by mystery shoppers during September 2015. The variable “Age Vendor” was not directly asked but guessed by mystery shoppers. The type of weight used to weight fish is a proxy of the level of formality of the fish stall.

The table C.3.2 describes the availability of different types of fish in feria stalls during the ban period. During a typical month, hake would be available in roughly 90% of stalls, however, due to the ban period (and our interventions), only 26% of stalls had hake for sale. The fish species offered in markets depend largely on the latitude of the market, i.e., markets located in the southern regions offer slightly different fish species than the stalls located in northern regions.

Table C.3.2: Fish Availability in Feria Stalls

Fish	Availability	Price/unit (USD)	Price/kg (USD)	Unitary Weight (kg)	N
Hake	0.263	1.08	3.81	0.284	239
Pomfret	0.684	5.27	4.99	1.057	621
Mackerel	0.124	2.05	4.16	0.492	113
Silverside	0.096	0.17	2.92	0.059	87
Salmon	0.139	7.53	9.22	0.816	126
Sawfish	0.057	6.27	6.25	1.003	52
Albacore	0.051	.	9.21	.	46
Southern Hake	0.042	7.30	5.59	1.306	38

Notes: This table presents the availability and average prices of different fish types in feria stalls during September 2015. The mystery shoppers recorded the price for each fish offered for sale in each fish-stall visited. The sale price in each stall was based on units, kilos or both. The unitary weight is estimated using the ratio of these two prices. The albacore is a considerably larger fish type (over 20 kgs) and is only sold in pieces (by kg).

Data Collected in the Fisherman Survey

A round of surveys to Fishermen was collected in August 2016. In total, 231 fishermen were surveyed and asked about their work, typical buyers and fishing behavior. The table C.3.3 describes the main variables collected in the survey.

Table C.3.3: Fishermen Characteristics

Variable	Mean	SD	Min	Max	N
Fisherman Boat characteristics:					
Boat Length (mts)	8.52	3.11	6	24	227
Boat Powered by a Motor	0.88	0.33	0	1	231
Fiberglass Boat	0.57	0.50	0	1	228
Wooden Boat	0.39	0.49	0	1	228
Union Participation:					
Number of Unions in the Caleta	1.67	0.90	0	3	230
Fisherman Member of a Union	0.82	0.38	0	1	230
Number of Days that Goes Fishing Every Week:					
Summer	5.01	1.47	1	7	226
Winter	2.25	1.14	0	7	227
Number of Boats in the Caleta:					
Less than 10	0.24	0.43	0	1	189
Between 10 and 30	0.22	0.41	0	1	189
Between 31 and 60	0.24	0.43	0	1	189
Between 61 and 100	0.12	0.32	0	1	189
More than 100	0.19	0.39	0	1	189
Top 3 Most Captured Fishes in the Caleta:					
Hake	0.56	0.50	0	1	230
Sawfish	0.24	0.43	0	1	230
Cuttlefish	0.24	0.43	0	1	230
Pomfret	0.13	0.34	0	1	230
Bass	0.10	0.30	0	1	230
Usual Buyer of the Fish at the Dock:					
Final Consumer	0.58	0.49	0	1	230
Feria Vendor	0.27	0.45	0	1	228
Intermediary	0.60	0.49	0	1	227

Notes: This table describes the responses to the Fishermen Survey carried out in August 2016 to 231 fishermen. On average, three fishermen were surveyed in each of the 74 caletas that operate in the four coastal regions included in our sample. The last section of the table represents the proportion that responded that *Always* or *Most of the Time* the fish was sold to these type of buyers.

Consumer Mobility Between Neighborhoods

The table C.3.4 shows the proportion of consumers treated by the information campaign depending on the location of the feria where they are surveyed. The striking fact in this table is that in high-saturation municipalities, the proportion of consumers treated with the information campaign is high, regardless of whether we found that person shopping in a feria located in a treatment neighborhood (78%) or in a control neighborhood (69%). High Information campaign saturation is therefore the effective treatment variable, and conditional on that, the specific location of the feria does not matter too much.

Table C.3.4: Proportion of Consumers located in Treated Neighborhoods

	Survey in Feria located in Treated Neigh		Survey in Feria located in Control Neigh	
	Prop	N	Prop	N
High Saturation Municipality	0.78	1114	0.69	389
Low Saturation Municipality	0.57	559	0.17	825
Zero Saturation Municipality	0	0	0.00	1014
Overall	0.71	1673	0.18	2228

Notes: This table shows the proportion of consumers whose households are located in neighborhoods assigned to receive the information campaign. These statistics are based on households' location reported by surveyed consumers. 71% of consumers surveyed in a feria located in a treated neighborhood live in a household located in a treated neighborhood; the remaining 29% are consumers who live in a household located in a control neighborhood. This table informs about the high consumers' mobility between neighborhoods. In fact, in high-saturation municipalities, the proportion of treated consumers is higher in both, ferias located in treatment and control neighborhoods.

Enforcement Implementation

This section describes the implementation of enforcement activities by Sernapesca officials. The research team planned the schedule of visits to different circuits. The execution was carried out by Sernapesca inspectors, as part of their usual tasks. The information about the actual visits was collected from the reports written by inspectors on a daily basis.[1] In total, 230 visits were carried out, equivalent to 659 stall-inspections in 62 circuits. Based on the inspectors' reports, illegal hake was detected in 11% of inspected stalls. This number is three times smaller than what our secret shoppers observed in markets.

Table C.3.5 describes the implementation of enforcement visits relative to the treatment assignment. The average number of visits in different treatment arms is slightly smaller than the original plan, this gap is explained by "contingencies" that obstructed the expected routine, and possibly, some under-reporting on behalf of inspectors. Also, a few visits were noted in Control group markets; these were generally markets located near Sernapesca regional offices, that officials unpromptedly visited.

[1] These reports contain information on the identity of the inspectors, the ferias visited that day, the number of fish stalls inspected, and whether illegal fish were detected. Importantly, the inspectors' performance does not depend on the information collected by these reports, but they rather work as a logbook of their activities. The research team periodically accessed, systematized and digitized this information.

Table C.3.5: Implementation of the Interventions

	(1)	(2)	(3)	(4)	(5)	(6)
				At Least One Visit		
Treatment Assignment	Number of Visits	Number of Different Days of the Week Visited	Circuit Visited at Least Once	Number of Visits	Number of Different Days of the Week Visited	N
No Enforcement	0.39	0.30	0.22	1.80	1.40	23
Enforcement	2.53	1.49	0.69	3.68	2.18	83
Low Freq. and Unpredictable Sch.	1.48	1.14	0.62	2.39	1.83	29
High Freq. and Unpredictable Sch.	5.00	2.80	0.87	5.77	3.23	15
Low Freq. and Predictable Sch.	1.30	0.85	0.65	2.00	1.31	20
High Freq. and Predictable Sch.	3.47	1.68	0.68	5.08	2.46	19

This table reports the unconditional mean of visits to circuits in different treatment arms. Column 4 presents the average number of visits conditional on receiving at least one visit. The difference between the number of visits in circuits assigned to High-Intensity Enforcement relative to Low Intensity is statistically significant at 1%.

The experimental design varied two margins of enforcement deployment; the frequency, and the predictability of the visits. The implementation of frequency variations can be evaluated based on the number of visits to each circuit. The predictability variation can be assessed based on the number of different days of the week in which visits were carried out. If predictable enforcement was implemented correctly, it should repeat the days of visits every week, so we should expect fewer days of the week visited. On average, circuits assigned to enforcement received 2.53 visits in 1.49 different days of the week. Both treatment variations generate significant differences in the relevant margins: High-frequency circuits received substantially more visits than low-frequency circuits. Unpredictable circuits were visited in more days of the week than predictable enforcement. The columns 4 and 5 compare these margins conditioning on receiving at least one visit.

Balance

We did not conduct a full baseline survey, but had access to municipality administrative data and weather data with which we could check balance across treatment arms. The table C.3.6 shows balance tests across the main treatment arms. Tables C.3.7 and C.3.8 also show balance tests with respect to the enforcement predictability and frequency sub-treatments.

Overall, the various treatment arms appear well balanced in terms of important socio-economic and weather characteristics (e.g. poverty rate, rainfall). The joint test F-statistics of all variables are insignificant for different treatment arms. The delinquency rate (i.e., per-capita police cases for major offenses) is lower in municipalities assigned to receive the information campaign relative to the control group. The regressions reported below control for this variable, but we have verified that the reported treatment effects are not sensitive to adding this control.

Balance Tables

The table C.3.6 presents the balance of relevant characteristics across different treatment arms. These variables include market's observable characteristics, socioeconomic characteristics of the municipality and weather information of the day of the visit by a mystery shopper. The columns 1, 2, 4 and 6 present the mean and SD of these variables for different treatment arms. The columns 3, 5 and 7 compare the difference relative to the control group as well as its p-value. Finally, joint significance tests are also reported in the last two columns. The tables C.3.7 and C.3.8 present the same estimates but decomposing by the enforcement variations: predictability and frequency.

Table C.3.6: Randomization Balance on Observables

Variable	(1) Control	(2) Info Campaign	(3)	(4) Enforcement	(5)	(6) Enforc. and Info Camp.	(7)
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573	0.644	0.083	0.489	-0.069	0.509	0.013
	(0.497)	(0.484)	[0.740]	(0.501)	[0.574]	(0.501)	[0.920]
Distance to Closest Caleta (kms)	16.507	11.572	3.465	14.863	-3.606	26.425	2.245
	(25.082)	(9.503)	[0.516]	(22.626)	[0.386]	(29.037)	[0.682]
Poverty Rate Municipality	19.006	17.567	-2.148	18.026	-1.079	16.734	-0.316
	(4.780)	(3.313)	[0.244]	(5.483)	[0.412]	(7.549)	[0.851]
Av. Monthly Income Municipality (USD)	791.767	875.475	18.446	790.514	-2.506	830.683	20.464
	(149.808)	(172.858)	[0.846]	(140.334)	[0.953]	(139.251)	[0.673]
Deliquency Rate Municipality	0.038	0.029	-0.013	0.036	-0.001	0.034	-0.004
	(0.015)	(0.002)	[0.016]	(0.015)	[0.835]	(0.009)	[0.480]
Rain Indicator	0.290	0.133	-0.124	0.178	-0.114	0.142	-0.122
	(0.456)	(0.344)	[0.455]	(0.383)	[0.318]	(0.349)	[0.301]
Average Temperature (Celsius)	12.200	12.126	0.081	11.993	-0.192	11.936	-0.688
	(2.281)	(2.087)	[0.942]	(2.021)	[0.797]	(2.196)	[0.346]
Joint Significance							
F statistic			0.609		1.094		0.816
p-value			0.747		0.371		0.575

Notes: This table reports characteristics of circuits included in our sample across treatment arms. The columns (1), (2), (4) and (6) show the mean and the standard deviation for the control and treatment groups. The columns (3), (5) and (7) show the coefficient on treatments from regressions of each characteristic on treatments and strata fixed effects, clustering standard errors at the circuit level. The p-values are reported in brackets. The socio-economic characteristics are aggregated at Municipality level. These variables should be interpreted as the characteristics of the Municipality where the circuit is located. Also, this table reports weather information of the day that different circuits were visited by mystery shoppers. Finally, joint significance test statistics: F statistic and p-values, for all variables on each treatment arm are reported in the last two rows of the table.

Table C.3.7: Randomization Balance: Enforcement Predictability

Variable	(1) Control	(2) Info Campaign	(3)	(4) Enforc: Predictable Schedule	(5)	(6) Enforc: Unpredictable Schedule	(7)
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573	0.644	0.080	0.396	-0.148	0.575	0.045
	(0.497)	(0.484)	[0.749]	(0.490)	[0.270]	(0.495)	[0.710]
Distance to Closest Caleta (kms)	16.507	11.572	3.822	17.286	1.704	20.487	-4.926
	(25.082)	(9.503)	[0.479]	(23.417)	[0.678]	(27.384)	[0.333]
Poverty Rate Municipality	19.006	17.567	-2.173	17.130	-1.739	17.890	-0.075
	(4.780)	(3.313)	[0.240]	(6.185)	[0.212]	(6.450)	[0.960]
Av. Monthly Income Municipality (USD)	791.767	875.475	18.828	809.457	-1.068	801.805	9.235
	(149.808)	(172.858)	[0.843]	(153.747)	[0.981]	(130.506)	[0.834]
Deliquency Rate Municipality	0.038	0.029	-0.013	0.036	-0.002	0.035	-0.001
	(0.015)	(0.002)	[0.015]	(0.015)	[0.648]	(0.012)	[0.782]
Rain Indicator	0.290	0.133	-0.126	0.117	-0.158	0.202	-0.080
	(0.456)	(0.344)	[0.446]	(0.322)	[0.162]	(0.402)	[0.488]
Average Temperature (Celsius)	12.200	12.126	0.076	12.058	-0.170	11.904	-0.491
	(2.281)	(2.087)	[0.946]	(2.057)	[0.824]	(2.107)	[0.498]
Joint Significance							
F statistic			0.609		1.954		1.717
p-value			0.747		0.067		0.111

Notes: This table reports characteristics of circuits included in our sample across treatment arms. The columns (1), (2), (4) and (6) show the mean and the standard deviation for the control and treatment groups. The columns (3), (5) and (7) show the coefficient on treatments from regressions of each characteristic on treatments and strata fixed effects, clustering standard errors at the circuit level. The p-values are reported in brackets. The socio-economic characteristics are aggregated at Municipality level. These variables should be interpreted as the characteristics of the Municipality where the circuit is located. Also, this table reports weather information of the day that different circuits were visited by mystery shoppers. Finally, joint significance test statistics: F statistic and p-values, for all variables on each treatment arm are reported in the last two rows of the table.

Table C.3.8: Randomization Balance: Enforcement Frequency

Variable	(1) Control	(2) Info Campaign	(3)	(4) Enforc: High Frequency	(5)	(6) Enforc: Low Frequency	(7)
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573	0.644	0.098	0.372	-0.164	0.591	0.039
	(0.497)	(0.484)	[0.696]	(0.484)	[0.211]	(0.492)	[0.753]
Distance to Closest Caleta (kms)	16.507	11.572	4.009	16.224	-5.302	21.269	0.600
	(25.082)	(9.503)	[0.451]	(23.309)	[0.330]	(27.292)	[0.885]
Poverty Rate Municipality	19.006	17.567	-1.975	16.242	-2.256	18.564	0.140
	(4.780)	(3.313)	[0.279]	(7.011)	[0.101]	(5.576)	[0.922]
Av. Monthly Income Municipality (USD)	791.767	875.475	17.128	821.315	23.237	792.779	-8.734
	(149.808)	(172.858)	[0.857]	(142.660)	[0.587]	(138.930)	[0.840]
Deliquency Rate Municipality	0.038	0.029	-0.013	0.036	-0.000	0.035	-0.003
	(0.015)	(0.002)	[0.014]	(0.013)	[0.973]	(0.014)	[0.547]
Rain Indicator	0.290	0.133	-0.128	0.193	-0.077	0.143	-0.143
	(0.456)	(0.344)	[0.439]	(0.396)	[0.491]	(0.350)	[0.215]
Average Temperature (Celsius)	12.200	12.126	0.050	12.127	-0.192	11.853	-0.447
	(2.281)	(2.087)	[0.964]	(2.204)	[0.799]	(1.984)	[0.545]
Joint Significance							
F statistic			0.609		2.016		2.090
p-value			0.747		0.058		0.049

Notes: This table reports characteristics of circuits included in our sample across treatment arms. The columns (1), (2), (4) and (6) show the mean and the standard deviation for the control and treatment groups. The columns (3), (5) and (7) show the coefficient on treatments from regressions of each characteristic on treatments and strata fixed effects, clustering standard errors at the circuit level. The p-values are reported in brackets. The socio-economic characteristics are aggregated at Municipality level. These variables should be interpreted as the characteristics of the Municipality where the circuit is located. Also, this table reports weather information of the day that different circuits were visited by mystery shoppers. Finally, joint significance test statistics: F statistic and p-values, for all variables on each treatment arm are reported in the last two rows of the table.

Table C.3.9 shows the coefficients of regression 3.2. The first three rows show the differences in the sale of illegal hake in the pre-intervention period. The interaction of "× Post" capture the effect that result from the interventions. The first three rows indicate that there were no statistically significant differences between treatment and control groups during the pre-intervention period. As expected, significant differences between markets appear after the interventions are launched (after the first week of September).

Table C.3.9: Treatment Effects on Hake Sales

VARIABLES	(1) Fresh, Visible Hake	(2) Any Hake Available (Hidden, Frozen, Visible)
Information Campaign Only	0.080	0.029
	(0.056)	(0.058)
Enforcement Only	0.114	0.092
	(0.070)	(0.060)
Information Campaign and Enforcement	0.078	0.100
	(0.070)	(0.065)
Information Campaign Only × Post	-0.133	-0.131
	(0.066)	(0.074)
Enforcement Only × Post	-0.178	-0.130
	(0.082)	(0.089)
Info Campaign and Enforcement × Post	-0.179	-0.139
	(0.074)	(0.094)
Change in Dep. Var. in Control Group During Intervention Period	-0.21	-0.36
N	901	901

Notes: This table reports the effect of each treatment arm on the availability of illegal hake fish. The variable Fresh Hake indicates when the hake was available fresh. Hake available indicates when was possible to buy fish in any form. The table reports marginal effects from a Probit regression. Other controls are included: municipality characteristics, strata fixed effects and the average level of the outcome variable in pre-intervention period. We control for pre-treatment values for the outcome variables in addition to the treatment indicator, because not all markets were visited in pre-intervention period. Robust standard errors clustered by circuit (the unit of randomization) in parentheses.

Adaptation to the Schedule of Visits

Our model suggests that vendors would learn and adapt to the pattern of visits. We use the daily data over the course of the September ban to study how the vendors adjust to the visit patterns they observe.

Table C.3.10 shows how selling decisions differ in the second half of the month, depending on how concentrated the earlier inspections were in specific ferias and on specific days of the week (DOWs). We control for the total number of visits in studying the effects of "targeting" only one feria or day-of-week. We find that auditing *different* ferias on *different* DOW reduces hake sales by an extra 9 percentage points (p-val¡0.01) in the second half of the month, relative to targeting all visits at the same feria.

Table C.3.11 studies vendors' decisions to sell hake in the *non-targeted* feria in the second half of the month. This is a circuit-fixed effects regression, so the coefficient "DOW not visited × 2nd half" shows the same vendor's decision to sell on a weekday in which he did not experience a visit relative to another weekday when he did. We see that the hake selling in the second half of the month was higher in ferias and DOWs that did *not* receive enforcement relative to the ones that did.

Table C.3.10: Hake Available based on the Number of different ferias and Days of the Week visited

	(1)	(2)	(3)
VARIABLES	Any Hake Available		
N Ferias Visited	0.041		0.039
	(0.030)		(0.033)
N Ferias Visited × Second Half	-0.091		-0.081
	(0.023)		(0.035)
N DOWs Visited		0.030	0.014
		(0.055)	(0.056)
N DOWs Visited × Second Half		-0.098	-0.037
		(0.073)	(0.077)
Change Dep Var First - Second Half	-0.31	-0.31	-0.31
N	906	906	906

Notes: This table studies how the probability of selling hake depends on the number of different days of the week (DOWs) and the number of different ferias that a circuit got visited during the ban. The observations are divided between the first and second half of the month to retain enough statistical power; other pre-post decompositions produce similar results. The table presents OLS coefficients of the relevant variables. Since DOWs and N Ferias are positively correlated, the columns 1 and 2 run them separately. Column 3 includes both variables and interactions. Each regression controls for "Second Half", the total number of visits, and the interaction of the two variables. Also, they control for the dependent variable in the pre-intervention period, and strata fixed effects and municipality characteristics. Also, each regression controls for treatment assignment. Robust standard errors clustered by circuit in parentheses.

We consider this evidence as only suggestive and placed it in the appendix, because Sernapesca chose which feria to visit within each circuit partly based on logistical considerations, and this cannot be treated as random. Indeed, the un-interacted terms in the regression show some differences (in the opposite direction!) across ferias within the circuit in the first part of the month.

Table C.3.12 shows that this same effect is not only seen in the propensity to sell, but also in the number of stalls that vendors choose to continue to operate in the second half of the month. The interacted coefficients "...× second half" show that more stalls disappear entirely in the second half of the month in the targeted ferias operating on targeted DOWs, relative to the non-targeted. The effect is larger when the vendors operate in more than two ferias, because those are the circuits where vendors have more options to adjust and displace sales across days of week.

The preceding tables explain why predictable enforcement is less effective. As our theoretical model lays out clearly, Vendors learn from the pattern of targeted ferias and targeted days of week, and adjust to sell more on non-targeted days.

Table C.3.11: Hake Sale across DOWs and Ferias within Circuit

	(1)	(2)	(3)	(4)
	Any Hake Available			
			Circuits that Rotate between	
VARIABLES	Full Sample		More Than 2 Ferias	
Feria Not Visited	-0.209		-0.299	
	(0.157)		(0.182)	
Feria Not Visited × Second Half	0.058		0.267	
	(0.086)		(0.130)	
DOW Not Visited		-0.252		-0.303
		(0.131)		(0.161)
DOW Not Visited × Second Half		0.192		0.286
		(0.104)		(0.139)
Change Dep Var First - Second Half	-0.31	-0.31	-0.39	-0.39
N	906	906	218	218

Notes: This table examines whether the behavior of the vendors varied across days of the week or ferias. It shows the OLS coefficient of dummy variables indicating whether the observation was collected in a feria or day that was not visited by Sernapesca officials during the ban. The observations are divided between the first and second half of the month to retain enough statistical power. Other pre-post decompositions produce similar results. These regressions include circuit fixed effects, so the coefficients capture within circuit variation. The columns (3) and (4) restrict the analysis only to circuits that rotate between more than two ferias. It controls for "Second Half" and weather covariates. Robust standard errors clustered by circuit in parentheses.

Exit of Stalls Correction

Our main results are based on the information gathered by mystery shoppers from the operative stalls at the moment of the visit, which does not capture the fact that the "missing" stalls are not selling hake. To correct for this issue we identify the average number of stalls per circuit/visit before and after the interventions. The comparison between these two averages informs about the number of "missing" stalls per circuit.[2] The number of stalls observed by mystery shoppers in every visit in the post treatment period is increased by computed number of missing stalls. The added observations have zero fish available.[3] [4]

[2] We allow the number of missing stalls to be non-integer, and negative if the number of stalls increased.

[3] If the number "missing" stalls is negative: the number of stalls observed by mystery shoppers in every visit in the pre-treatment period is increased by that number.

[4] Since we allow the number "missing" stalls to be non-integer, we add a random noise that distributes uniform between -0.5 and 0.5, and then, the sum of the "missing" number and the noise is rounded to the closest integer. This correction makes the expansion more representative of the right (possibly non-integer) number.

Table C.3.12: Number of Stalls Hake Sale across DOWs and Ferias within Circuit

	(1)	(2)	(3)	(4)
	Number of Stalls			
VARIABLES	Full Sample		Circuits that Rotate between More Than 2 Ferias	
Feria Not Visited	-0.510		-0.956	
	(0.320)		(0.368)	
Feria Not Visited × Second Half	0.354		1.331	
	(0.464)		(0.771)	
DOW Not Visited		-0.024		-0.241
		(0.220)		(0.312)
DOW Not Visited × Second Half		0.317		0.995
		(0.330)		(0.450)
Change Dep Var First - Second Half	0.03	0.03	0.20	0.20
N	374	374	104	104

Notes: This table examines whether the number of stalls selling fish varied across days of the week or ferias. It shows the OLS coefficient of dummy variables indicating whether the observation was collected in a feria or day that was not visited by Sernapesca officials during the ban. Every observation correspond to a feria and are divided between the first and second half of the month to retain enough statistical power. Other pre-post decompositions produce similar results. These regressions include circuit fixed effects, so the coefficients capture within circuit variation. The columns (3) and (4) restrict the analysis only to circuits that rotate between more than two ferias. It controls for "Second Half" and weather covariates. Robust standard errors clustered by circuit in parentheses.

Table C.3.13: Treatment Effect on Hake Availability Correcting for the Exit of Stalls

VARIABLES	(1) Fresh, Visible Hake	(2) Any Hake Available (Fresh-Visible, Hidden or Frozen)
Panel A: Main Specification		
Info Campaign Only	-0.118	-0.115
	(0.060)	(0.065)
Enforcement Only	-0.190	-0.141
	(0.082)	(0.091)
Info Campaign and Enforcement	-0.156	-0.130
	(0.084)	(0.104)
Panel B: Variation in Predictability of Enforcement		
Info Campaign Only	-0.111	-0.121
	(0.062)	(0.064)
Enforcement on Predictable Schedule	-0.091	-0.061
	(0.073)	(0.087)
Enforcement on Unpredictable Schedule	-0.246	-0.197
	(0.089)	(0.100)
Panel C: Variation in Frequency of Enforcement		
Info Campaign Only	-0.113	-0.121
	(0.062)	(0.064)
High Frequency Enforcement	-0.086	-0.092
	(0.092)	(0.101)
Low Frequency Enforcement	-0.184	-0.148
	(0.086)	(0.095)
Change in Dep Var in Control		
During Intervention	-0.17	-0.28
Covariates	Yes	Yes
Baseline Control	Yes	Yes
N	1014	1014

Notes: This table presents the coefficient corresponding to the interaction term $T_c \times Post_t$ for each treatment correcting for the exit of stalls. The increase in the number of observations relative to results presented earlier is due to the fact that the correction is done by adding the "missing" stalls (calculated compraing the number of stalls per circuit before and after the interventions). The panel A describes the same specification presented in Table 3.2. Panels B and C follow the same specification than Table 3.4. Probit regression marginal effects are reported. Robust standard errors clustered by circuit in parentheses.

Treatment Effects Six Months After the Ban Period

Table C.3.14 describe the answers to the consumer survey carried out in March 2016. The survey had the same format as previous surveys; it asked about general consumption behavior based on a list of items, including hake. Even though the survey was carried out in an off-ban period, consumers assigned to the information campaign tend to report less hake consumption.

Table C.3.14: Hake Purchases Reported by Consumers in March 2016 (Outside Ban Period)

VARIABLES	(1) Purchased Hake last month	(2) Number Times Hake Purchased	(3) Usually Purchase Hake
Information Campaign Only	-0.133	-0.419	-0.121
	(0.106)	(0.234)	(0.100)
Enforcement Only	0.018	-0.106	0.093
	(0.078)	(0.205)	(0.069)
Info Campaign and Enforcement	-0.017	-0.197	0.017
	(0.081)	(0.246)	(0.077)
Mean Dep Var Control	0.59	1.19	0.58
N	3652	3630	3652

Notes: This table presents the effect of different treatment arms on the reported consumption of hake fish by consumers based on the round of surveys collected in March 2016. Columns 1 and 3 show marginal effects from Probit regressions. Column 2 shows the marginal effects from a Poisson regression because the dependent variable is a count data. These regressions control for propensity to purchase other types of fishes and other covariates. Standard errors are clustered based on the circuit where the survey was collected.

Alternative Definition Information Campaign Treatment

Tables C.3.15 and C.3.16 present the main results using a different definition of the Information Campaign treatment: The variable "Information campaign" indicates whether the observations were collected by mystery shoppers in ferias located in treated neighborhoods - regardless of the level of saturation of the municipality. This definition does not include possible information spill-overs between neighborhoods within municipalities assigned to receive information.

Table C.3.15: Treatment Effect on Hake Availability

VARIABLES	(1) Fresh, Visible Hake	(2) Any Hake Available (Fresh-Visible, Hidden or Frozen)
Information Campaign Only	-0.082	-0.070
	(0.064)	(0.071)
Enforcement Only	-0.157	-0.101
	(0.079)	(0.094)
Info Campaign and Enforcement	-0.169	-0.121
	(0.079)	(0.094)
Change in Dep Var in Control Markets During Intervention	-0.21	-0.36
N	901	901

Notes: This table reports the effect of each treatment arm on the availability of illegal hake fish. The variable "Info Campaign" indicates if the feria where the data was collected is located in a neighborhood assigned to receive the information campaign. Probit Marginal effects of he interactions $T_c \times Post_t$ are reported. Robust standard errors are clustered by circuit and presented in parentheses.

Table C.3.16: Treatment Effect on Hake Sales by Enforcement Strategy

VARIABLES	(1) Any Hake	(2) Available
Info Campaign Only	-0.073	-0.073
	(0.071)	(0.071)
Enforcement on Predictable Schedule	-0.036	
	(0.089)	
Enforcement on Unpredictable Schedule	-0.169	
	(0.099)	
High Freq. Enforcement		-0.049
		(0.101)
Low Freq Enforcement		-0.140
		(0.095)
Change in Dep Var in control Markets During Intervention	-0.36	-0.36
N	901	901

Notes: This table reports the effect of each treatment arm on the availability of illegal hake fish. The first column includes compares the effectiveness of predictable vs unpredictable enforcement. The second column divides enforcement depending on its intensity. Each regression controls for the dependent variable in pre-intervention period, strata fixed effects and municipality characteristics. The variable "Info Campaign" indicates if the feria where the data was collected is located in a neighborhood assigned to receive the information campaign. Probit Marginal effects of he interactions $T_c \times Post_t$ are reported. Robust standard errors are clustered by circuit and presented in parentheses.

C.4 Appendix: Theoretical Model

Belief Formation with More than One Feria

We denote by $z_t = y_t^2 + 2y_t^1 + 1$ the multinomial random variable of the profile of inspections in period t, which we assumes is the underlying distribution determining the probability if a visit in each feria.[5] We assume that z_t has a stationary distribution characterized by $p = (p^j)_{j=1}^4$, where $p^j = \mathbb{P}(z_t = j)$. In this case we denote the prior by $\hat{p}_0 \sim \text{Dirichlet}\left((\beta_i)_{i=1}^4\right)$. Finally, we denote by $\theta = (\theta^1, \theta^2)$ the real probability of visits, which we call the *visit policy*.[6] Note that $\theta^1 = p^2 + p^4$ and $\theta^2 = p^3 + p^4$.

The following result extends the vendor's belief dynamics for this case.

Lemma 1. *The vendor's belief about the probability of a visit at feria i at time t satisfies*

$$\hat{\theta}_t^1 \sim Beta\left(\alpha_2 + \alpha_4 + Y_t^1; \alpha_1 + \alpha_3 + t - 1 - Y_t^1\right)$$
$$\hat{\theta}_t^2 \sim Beta\left(\alpha_3 + \alpha_4 + Y_t^2; \alpha_1 + \alpha_2 + t - 1 - Y_t^2\right)$$

This result shows that the vendor updates her beliefs about the probability of a visit in each feria by looking only at the history of visits at that feria.

Proofs

Proof of Proposition 1. First we analyze the vendor's the different options. To avoid unnecessary notation we omit the subindex t and write $g = g(Y_t)$.

- She sells and does not defend if and only if $U[h = 0|s = 1, Y] \geq 0$ and $U[h = 0|s = 1, Y] \geq U[h = 1|s = 1, Y]$. These restrictions together imply that

$$\mathbb{E}[\hat{\theta}] \leq \frac{1}{\Omega} \min\left\{v; \frac{c}{g}\right\}.$$

- She sells and defends if and only if $U[h = 1|s = 1, Y] \geq 0$ and $U[h^i = 1|s^i = 1, Y] > U[h^i = 0|s^i = 1, Y]$. These restrictions together imply that (recall that $\underline{\delta} = \frac{c}{\Omega g}$, and $\overline{\delta} = \frac{v-c}{\Omega(1-g)}$)

$$\underline{\delta} < \mathbb{E}[\hat{\theta}] \leq \overline{\delta}.$$

- The vendor does not sell if and only if $\max_{h\in\{0,1\}} U[h|s = 1, Y] < 0$. These restrictions together imply that

$$\mathbb{E}[\hat{\theta}] > \frac{1}{\Omega} \max\left\{v; \frac{v-c}{1-g}\right\}.$$

[5] Note that z_t takes value one if no feria was inspected at time t, value 2 if only feria 1 was inspected, value 3 if only feria 2 was inspected, and value 4 if both ferias were inspected in that period.

[6] As the distribution of z_t is stationary, the probabilities of visits in both ferias also are.

First, note that the conditions $\underline{\delta} < \overline{\delta}$, $v > \frac{c}{g}$, and $v < \frac{v-c}{1-g}$ are equivalent. Therefore, there is a set of beliefs for which the vendor's optimal choice is to sell and defend the hake if and only if $v > \frac{c}{g}$.

If $v \leq \frac{c}{g}$, then $\underline{\delta} \geq \overline{\delta}$ and the vendor never sells and defends. Moreover, as $\min\left\{v; \frac{c}{g}\right\} = v$ in this case she sells and does not defend if $\mathbb{E}[\hat{\theta}] \leq \frac{v}{\Omega}$ and does not sell otherwise.

Finally, if $v \leq \frac{c}{g}$ the characterization follows directly form the comparison of the three options. □

Long-run Comparative Statics. Define $\underline{\delta}_\infty = \frac{c}{\Omega \bar{g}}$ and $\overline{\delta}_\infty = \frac{v-c}{\Omega(1-g)}$. First, note that Assumption $\bar{g} > cv$ implies that $\underline{\delta}_\infty < \overline{\delta}_\infty$. As (a.s.) $\underline{\delta}_t \to \underline{\delta}_\infty$ and $\overline{\delta}_t \to \overline{\delta}_\infty$, we have that in the long run there is a set of beliefs for which the vendor sells and defends. Furthermore, Proposition 1 implies that the vendor sells if and only if

$$\theta \leq \overline{\delta}_\infty = \frac{v-c}{\Omega(1-\bar{g})} .$$

The comparative statics results follow from analyzing the effect of changes in θ, v, and $1-\bar{g}$ in the previous inequality. □

Proof of Lemma 1. For any t we define the number of periods *before* t that the vendor has seen $z = j$ (for $j = 1, 2, 3, 4$) by

$$Z_t^j = \sum_{s=1}^{t-1} \mathbb{1}_{\{z_s = j\}}$$

Given $Z_t = \left(Z_t^j\right)_{j=1}^4$, Bayesian updating implies that

$$\hat{p}_t \sim \mathrm{Dir}\left(\alpha + Z_t\right)$$

As $Z_t^2 + Z_t^4 = Y_t^1$, the previous distribution implies that the vector

$$\left(\hat{p}_t^1, \hat{p}_t^2 + \hat{p}_t^4, \hat{p}_t^3\right) = \left(\hat{p}_t^1, \hat{\theta}_t^1, \hat{p}_t^3\right) \sim \mathrm{Dir}\left(\alpha^1 + Z_t^1, \alpha^2 + \alpha^4 + Y_t^1, \alpha^3 + x_t^3\right).$$

As $\sum_{j=1}^4 \hat{p}_t^j = 1$ and $\sum_{j=1}^4 Z_t^j = t-1$, the marginal distribution of the probability of being inspected at feria 1 at time $t+1$ is

$$\hat{\theta}_t^1 \sim \mathrm{Beta}\left(\alpha^2 + \alpha^4 + Y_t^1, \alpha^1 + \alpha^3 + t - 1 - Y_t^1\right).$$

The characterization of the distribution of $\hat{\theta}_t^2$ is completely analogous. □

Proof of Proposition 2. First, note that it is a direct extension of Proposition 1 to show that in the long run the vendor sells in feria i if and only if $\theta^i \leq \frac{v-c}{\Omega(1-\bar{g})}$.

To show that it is without loss of generality to focus only on targeted and unpredictable policies, take any policy (θ^1, θ^2) such that $\theta^1 + \theta^2 = \Theta$.

- If the policy does not prevent selling in any feria, it is clear that both the targeted and the unpredictable policies are weakly more efficient.
- If the policy prevents selling only in feria i, we have that $\theta^i > \frac{v-c}{\Omega(1-\bar{g})} \geq \theta^{-i}$. As $\theta^i \leq \Theta$, we have that the targeted policy targeting feria 1 (or feria 2) is weakly more efficient.
- If the policy prevents selling in both ferias, we have that $\theta^1, \theta^2 > \frac{v-c}{\Omega(1-\bar{g})}$. As $\Theta 2 \geq \min\{\theta^1; \theta^2\}$, we have that the unpredictable policy $(\Theta 2, \Theta 2)$ also prevents selling in both ferias.

Now we analyze the most efficient policy for different values of Θ:

1. If $\Theta < \frac{v-c}{\Omega(1-\bar{g})}$: In this case neither the targeted policy or the unpredictable policy prevent selling in any feria.
2. If $\frac{v-c}{\Omega(1-\bar{g})} \leq \Theta < 2\frac{v-c}{\Omega(1-\bar{g})}$. In this case then the targeted policy targeting feria 1(2) prevents selling in feria 1(2) and does not prevent selling in feria 2(1). On the other hand, as $\Theta 2 < \frac{v-c}{\Omega(1-\bar{g})}$ the unpredictable policy does not prevent selling in any feria.
3. If $\Theta \geq 2\frac{v-c}{\Omega(1-\bar{g})}$ then the unpredictable policy prevents selling in both ferias, while the targeted policy prevents selling in only one of them.

□

The vendor's ability to circumvent the fine reaches a static value $\bar{g}$ in the long-run. So she only sells in a feria if her perceived probability of an enforcement visit is below the threshold $\overline{\delta}_t$. Hence, in the log-run, illegal selling is avoided in a feria if its inspection intensity θ^i is above the threshold. Furthermore, as the total enforcement capacity Θ is fixed, the policy can either reach the threshold in both ferias, in only one feria, or in neither feria. If enforcement capacity is not high enough to reach $\overline{\delta}_t$ in both ferias, the inspector should choose a targeted policy to prevent illegal sales in at least one feria

Proof of Corollary 2. We analyze the different cases characterized in Proposition 1. For the analysis we use that $\mathbb{E}[\hat{\theta}_t] \lesseqgtr \underline{\delta}_t \iff \Omega\mathbb{E}[\hat{\theta}_t] \lesseqgtr cg(Y_t)$, and $\mathbb{E}[\hat{\theta}_t] \lesseqgtr \overline{\delta}_t \iff \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c \lesseqgtr v$.

- The agent does not sell in two cases
    1. If $v \leq cg(Y_t)$, the vendor does not sell if $v < \Omega\mathbb{E}[\hat{\theta}_t]$. This happens with probability

$$\mathbb{P}\left(v \leq \frac{c}{g(Y_t)}\right)\mathbb{P}\left(v < \Omega\mathbb{E}[\hat{\theta}_t] \middle| v \leq \frac{c}{g(Y_t)}\right) = \mathbb{P}\left(\left(v < \Omega\mathbb{E}[\hat{\theta}_t]\right)\left(v \leq \frac{c}{g(Y_t)}\right)\right)$$
$$= \mathbb{P}\left(v < \min\left\{\frac{c}{g(Y_t)}; \Omega\mathbb{E}[\hat{\theta}_t]\right\}\right)$$
$$= F\left(\min\left\{\frac{c}{g(Y_t)}; \Omega\mathbb{E}[\hat{\theta}_t]\right\}\right).$$

2. If $v > cg(Y_t)$ the vendor does not sell if $\mathbb{E}[\hat{\theta}_t] > \overline{\delta}_t$.[7] The probability of this is

$$\mathbb{P}\left(v > \frac{c}{g(Y_t)}\right)\mathbb{P}\left(v < \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c \middle| v > \frac{c}{g(Y_t)}\right) =$$
$$= \mathbb{P}\left(\left(v < \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c\right) \left(v > \frac{c}{g(Y_t)}\right)\right)$$
$$= \mathbb{P}\left(\frac{c}{g(Y_t)} < v < \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c\right)$$
$$= \max\left\{F\left(\Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c\right) - F\left(\frac{c}{g(Y_t)}\right); 0\right\}.$$

The share of vendors that do not sell α_{NS} is the sum of these two probabilities.

- The vendor sells openly in two cases

1. If $v \le cg(Y_t)$, the vendor does not sell if $v > \Omega\mathbb{E}[\hat{\theta}_t]$. This happens with probability

$$\mathbb{P}\left(v \le \frac{c}{g(Y_t)}\right)\mathbb{P}\left(v > \Omega\mathbb{E}[\hat{\theta}_t] \middle| v \le \frac{c}{g(Y_t)}\right) = \mathbb{P}\left(\left(v > \Omega\mathbb{E}[\hat{\theta}_t]\right) \left(v \le \frac{c}{g(Y_t)}\right)\right)$$
$$= \mathbb{P}\left(\Omega\mathbb{E}[\hat{\theta}_t] < v < \frac{c}{g(Y_t)}\right)$$
$$= \max\left\{F\left(\frac{c}{g(Y_t)}\right) - F\left(\Omega\mathbb{E}[\hat{\theta}_t]\right); 0\right\}.$$

2. If $v > cg(Y_t)$ the vendor sells openly if $\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t$.[8] The probability of this is

$$\mathbb{P}\left(v > \frac{c}{g(Y_t)}\right)\mathbb{P}\left(\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t \middle| v > \frac{c}{g(Y_t)}\right) = \mathbb{P}\left(\left(\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t\right) \left(v > \frac{c}{g(Y_t)}\right)\right)$$
$$= \mathbb{P}\left(\frac{c}{g(Y_t)} < v\right)\mathbb{1}_{\{\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t\}}$$
$$= \left(1 - F\left(\frac{c}{g(Y_t)}\right)\right)\mathbb{1}_{\{\Omega\mathbb{E}[\hat{\theta}_t] \le \frac{c}{g(Y_t)}\}}.$$

The share of vendors who sell openly α_{SO} is the sum of these two probabilities.

- The vendor sells defensively only if $\overline{\delta}_t < \mathbb{E}[\hat{\theta}_t] \le \overline{\delta}_t$.[9] The share of vendors who sell

[7] If $v > cg$ and $\mathbb{E}[\hat{\theta}_t] > \overline{\delta}_t$, the condition $\mathbb{E}[\hat{\theta}_t] > \underline{\delta}_t$ is necessarily satisfied.
[8] If $v > cg$ and $\mathbb{E}[\hat{\theta}_t] \le \underline{\delta}_t$ the condition $\mathbb{E}[\hat{\theta}_t] \le \overline{\delta}_t$ is necessarily satisfied.
[9] Recall that the conditions $\underline{\delta}_t < \overline{\delta}_t$ and $v > cg(Y_t)$ are equivalent.

defensively α_{SD} is

$$\begin{aligned}\mathbb{P}\left(\left(\mathbb{E}[\hat{\theta}_t] > \underline{\delta}_t\right)\left(v \geq \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t))+c\right)\right) &= \mathbb{P}\left(v \geq \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t))+c\right)\mathbb{1}_{\{\mathbb{E}[\hat{\theta}_t]>\underline{\delta}_t\}}\\ &= \left(1-F\left(\Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t))+c\right)\right)\mathbb{1}_{\{\Omega\mathbb{E}[\hat{\theta}_t]>\frac{c}{g(Y_t)}\}}\end{aligned}$$

To finish the proof we just need to note that $\Omega\mathbb{E}[\hat{\theta}_t] \leq cg(Y_t) \iff \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) - c \leq cg(Y_t)$ and replace the corresponding values in each case. □

C.5 Numerical Simulations

We numerically simulate the behavior of a representative vendor exposed to different levels and schemes of enforcement. These simulations shed light on how the optimal choice evolves as vendors acquire more information about the pattern of visits and inspection loopholes.

Vendors' Behavior Over Time Vendors decide whether and how to sell hake in every period t. The decision to sell in t is static but incorporates the information collected until $t-1$. Thus it may vary as more information is incorporated. In particular, vendors continuously update their probability of receiving an enforcement visit as well as the effectiveness of defensive strategies reducing the probability of a fine. Figure C.5.1 describes how the optimal decision in different periods. It shows that the likelihood of selling is not stable. In this case, it decreases quickly once the enforcement is introduced, and increases as the vendor learn about enforcement weaknesses. After a few periods, it converges to the long-run equilibrium. One direct takeaway of C.5.1 is that vendors' behavior varies over time; the same policy evaluated in different moments may yield different results.

Figure C.5.1: Vendor's Decision

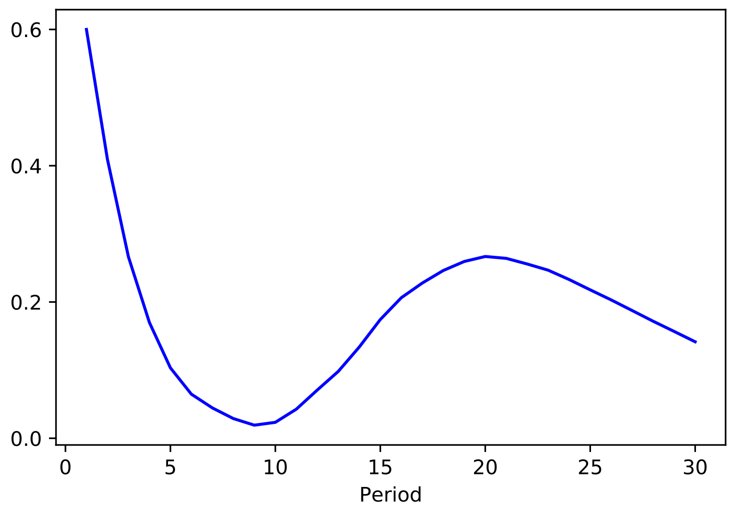

This figure shows the proportion of times in which a vendor sells hake in different periods. This graph depicts 1000 simulations using the following parameters $\theta = 0.4, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + e^{-8\times Y+28}\right)$, i.e.,$\bar{g} = 0.7$. The probability of selling decreases quickly as the enforcement begins, however it increases as vendors learn about enforcement weaknesses. After a number of periods, it converges to the "long-run" equilibrium based on model's structural parameters.

Enforcement Intensity Vendors adapt their behavior according to the pattern of visits they receive. We compare the behavior of vendors exposed to different frequencies of visits. Figure 3.2 shows that vendors exposed to more intense enforcement tend to decrease the probability of selling quickly. However, they learn faster about enforcement weaknesses. Thus, after a few periods, the latter effect may counterbalance the higher intensity effect, which makes high-intensity enforcement less effective. As the number of periods increases, the selling decision converges to the long-run optimal. This result has relevant implications for enforcement evaluation and design.

The figures C.5.2a and C.5.2b describe the timing and scope of adoption of defensive actions depending on the frequency of the enforcement. Vendors exposed to more intense enforcement learn quickly about loopholes, so they start adopting these actions earlier.

Figure C.5.2: Probability of Selling Hake

(a) Low Frequency Enforcement

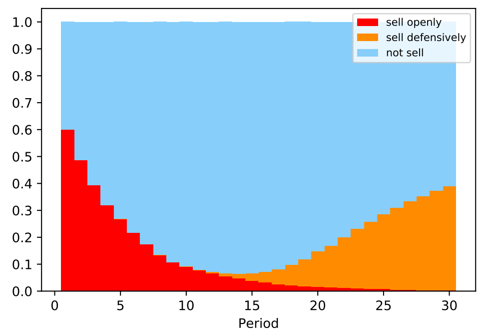

(b) High Frequency Enforcement

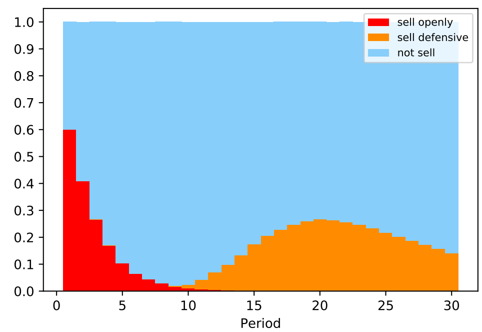

Figure C.5.2a and C.5.2b describe vendors' decision on whether and how to sell. This simulation uses the same parameters than previous graph: $\theta^{high} = 0.5, \theta^{low} = 0.3, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + e^{-2\times Y+12}\right)$, i.e.,$\bar{g} = 0.7$. The adoption of defensive strategies starts after a number of periods.

Enforcement Predictability We study the consequences of varying the predictability of the enforcement visits. In particular, we study vendors' behavior, assuming that every circuit has two ferias and that the vendor alternates between them. If enforcement is predictable, one of the ferias receives enforcement more intensely than the other. In our analysis, the probability of receiving a visit in a non-targeted feria is zero. Conversely, under unpredictable enforcement, both ferias have the same likelihood of receiving a visit.

The C.5.3a and C.5.3b show how, under predictable enforcement, the behavior of vendor diverge across ferias, the probability of selling in a non-targeted feria tend to one, whereas in a targeted feria tends to zero. i.e., the average tends to 0.5. The speed of convergence to 0.5 hinges on the overall enforcement frequency.

Figure C.5.3: Probability of Selling Hake

(a) Low Frequency Enforcement (b) High Frequency Enforcement

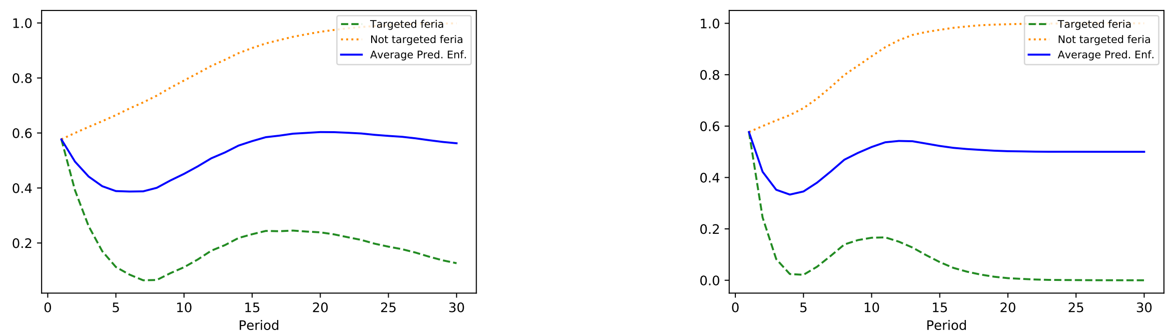

Figure C.5.3a and C.5.3b describe vendors' decision vary depending on the feria they are selling. The model assumes vendors alternate between taregeted and non-targeted feria. These simulations assume that in every period there's half of the vendors in each type of feria. The dashed line correspond to the average probability of sale. This simulation uses assumes $\theta^{high} = 0.4, \theta^{low} = 0.25, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + e^{-8\times Y+28}\right)$, i.e.,$\bar{g} = 0.7$.

The figures C.5.4a and C.5.4b compare the average probability of selling hake under predictable and unpredictable enforcement using the same enforcement capacity. As discussed in section 3.3, the long-run effects of one policy over the other depending on the structural parameters of the model. However, in the short-run, the speed and scope of learning play a role. Under most functional forms, the unpredictable enforcement seems to be more effective in the short-run. Illegal sales in ferias fall sharply as soon as auditors start visiting, but under predictable enforcement, the non-targeted feria does not benefit from this.

Figure C.5.4: Probability of Selling Hake

(a) Low Frequency Enforcement

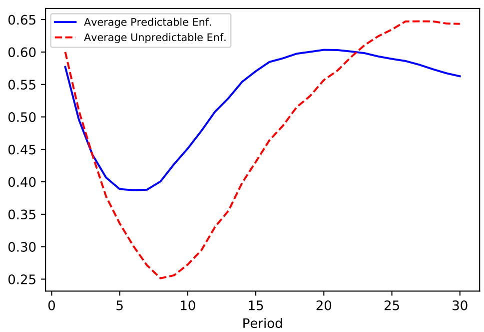

(b) High Frequency Enforcement

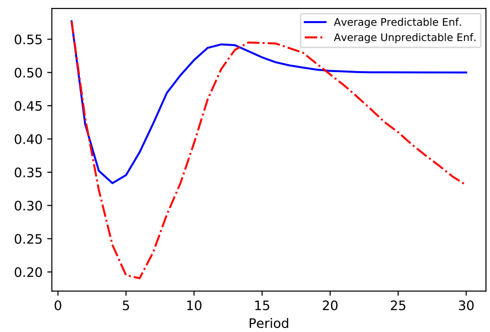

Figure C.5.4a and C.5.4b describe vendors' decision vary depending on the feria they are selling. The model assumes vendors alternate between taregeted and non-targeted feria. These simulations assume that in every period there's half of the vendors in each type of feria. The dashed line correspond to the average probability of sale. This simulation uses assumes $\theta^{high} = 0.4, \theta^{low} = 0.25, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7/\left(1 + e^{-8\times Y+28}\right)$, i.e.,$\bar{g} = 0.7$.

Note about Agents Heterogeneity We introduce agents' heterogeneity by agents that differ in their valuation v. Specifically, suppose v is distributed according to the CDF F, whose support is $[\underline{v}, \overline{v}]$. Assume $c < \underline{v} \leq \overline{v} < \Omega$. The applying 1 we get the following result

Corollary 2. *In the case with heterogeneous agents let α_{NS}, α_{SO}, and α_{SD} be the share of agents not selling, selling openly, and selling defensively, respectively. This shares are given by*

$$\alpha_{NS} = \begin{cases} F\left(\Omega\mathbb{E}[\hat{\theta}_t]\right) & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] \leq cg(Y_t) \\ F\left(\Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c\right) & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] > cg(Y_t); \end{cases}$$

$$\alpha_{SO} = \begin{cases} 1 - F\left(\Omega\mathbb{E}[\hat{\theta}_t]\right) & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] \leq cg(Y_t) \\ 0 & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] > cg(Y_t); \ \textit{and} \end{cases}$$

$$\alpha_{SD} = \begin{cases} 0 & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] \leq cg(Y_t) \\ 1 - F\left(\Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t)) + c\right) & \textit{if } \Omega\mathbb{E}[\hat{\theta}_t] > cg(Y_t). \end{cases}$$

Assumptions The numerical simulations presented above assume a functional form the the learning curve $g(Y_t)$ and a set (and strength) of priors. In particular, we assume: $g(Y_t) = \frac{\bar{g}}{1+exp\{-a\times Y+b\}}$. This functional form is handy because, $\lim_{x\to\infty} g(x) = \bar{g}$, and the parameters a and b dictate the speed of convergence of the function. Other functional forms yield the same qualitative results. The figures C.5.5a and C.5.5b describe how the ability and the beliefs evolve over time for two different levels of enforcement. The idea is that vendors exposed to more intense enforcement develop an ability to circumvent the fine faster, this effects counterbalances the increase probability of a visit.

Figure C.5.5: Beliefs Updating and Learnning Curve

(a) Ability

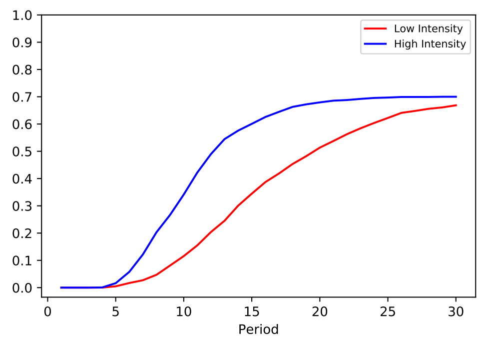

(b) Beliefs

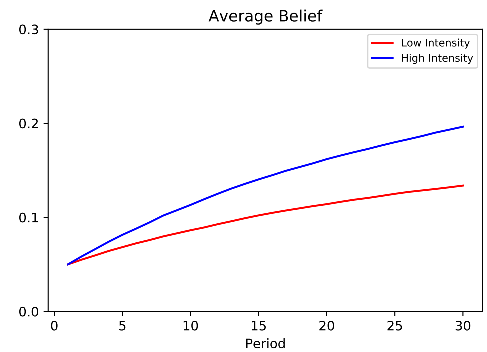

Figure C.5.5a and C.5.5b how vendors learn about loopholes and update the probability of a visit. These figures use the same inputs than other simulations, i.e., $\bar{g} = 0.7, a = 8, b = 28, \theta_1 = 0.05\,(\alpha_0 + \alpha_1 = 40)$

C.6 Further Details on the Cost-Effectiveness Analysis

The section 3.7 describes the cost-effectiveness analysis of each treatment. These calculations are based on the following parameters:

- Costs: The total cost of implementing enforcement was \$ 62,900.25, which is divided into fixed costs \$ 7,338.06, and variable costs: \$ 55,562.19. The fixed costs include administrative staff salaries, central office coordination, etc. The variable costs include the specific costs incurred to implement the enforcement (i.e., financial compensation of inspectors, gasoline, etc.). Based on Sernapesca information, deploying enforcement in an unpredictable way is 10% more costly regarding staff availability. The cost of implementing enforcement at low frequency is obtained by calculating the (variable) cost of each visit and multiplying by the number of visits under this new regime, adding the fixed costs.

  The total cost of implementing the information campaign was \$ 16,213.53, which includes the printing and distribution of flyers, posters, and letters in treated neighborhoods.

- Reduction of fish sales: The estimated effects of selling hake during the ban presented in section 3.5 are translated into numbers of fishes "saved." This exercise takes into account that every stall has 25 hake fishes available, there are 2.57 fish stalls in each feria. Each circuit operates 5 days a week, and the effects consider the three last weeks of September. The enforcement treatment contemplated 83 circuits, whereas 26 circuits are located in municipalities assigned to receive information campaign with a high level of saturation. The information gathered from the vendors surveyed provided useful information to define the right parameters regarding the likely reduction on fish sales as a result of our interventions.

C.7 Departures from the Pre-Analysis Plan

We registered this trial on September 15, 2015 (before the data collection was completed) in the AEA registry. Our approach to analysis and the outcome variables we focus on in this paper closely mirror the project narrative we uploaded before we had access to any data. We highlight the most notable departures from the pre-analysis plan (PAP) here:

1. The experimental design section of the PAP mentions that the enforcement group would be divided into two sub-groups: One in which vendors would receive only a warning letter about illegal behavior, and one in which we would follow that up with inspections and fines. In practice, *Sernapesca* officials did not implement the treatments any differently across these two sub-groups. So we do not report this sub-sample analysis. Our data show that vendor behavior was not statistically distinguishable across these sub-groups.

2. The PAP mentions our sample size as 153 circuits, based on information we had collected on the existence of fish markets by calling municipalities before launching the project. During data collection we learnt that 40 of those circuits did not have any fish-stalls. Mystery shoppers could not visit another 7 circuits for logistical reasons. Our final analysis sample therefore contains only 106 circuits. These two sources of attrition are not correlated with any observable characteristics, nor with the treatment assignment.

3. We had not anticipated that vendors would try to cheat by claiming that the fish was caught in August. This is something we learnt from our mystery shoppers soon after we started data collection. In the PAP, we mention only that we will track vendor reactions to enforcement activity, but do not mention 'freezing' specifically.

4. The PAP does not delve into the level of detail that this paper does. For example, we did not know exactly which fish were close substitutes for hake. We learnt from our data that pomfret was the other fish most commonly sold by hake vendors, and we therefore analyze effects on the price of pomfret. This price analysis could therefore

be viewed as "exploratory" even though we had pre-specified our interest in studying price effects.