

SLIPPERY FISH:
ENFORCING REGULATION WHEN AGENTS LEARN AND ADAPT

By

Andres Gonzalez-Lira and Ahmed Mushfiq Mobarak

August 2018

Revised March 2021

COWLES FOUNDATION DISCUSSION PAPER NO. 2143R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Slippery Fish:

Enforcing Regulation when Agents Learn and Adapt

Andres Gonzalez-Lira* Ahmed Mushfiq Mobarak†

March 7, 2021

Abstract

Attempts to curb undesired behavior through regulation gets complicated when agents can adapt to circumvent enforcement. We 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. 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.

JEL Codes: K42, O1, L51

Keywords: Enforcement, Regulation, Law and Economics, Fisheries.

*UC Berkeley, andres.gonzalez@berkeley.edu

†Yale University, Deakin University, CEPR and NBER, ahmed.mobarak@yale.edu

‡We thank the National Marine Resource Authority of Chile, *Sernapesca*, for their cooperation in implementing a randomized controlled trial; *JPAL-Latin America* for providing research and data collection support as well as the connection to Sernapesca through a JPAL Executive Education program in Santiago; the Macmillan Center at Yale, the Center for Business and Environment at Yale, and *Sernapesca* for financial support; Gunjan Amarnini, Olivia Bordeu, Pascuala Domínguez, Angélica Eguiguren, Matthew Sant-Miller, Cristian Ugarte and Diego Verdugo for research assistance; Anjali Adukia, Kaushik Basu, Eli Berman, Fiona Burlig, Christopher Costello, Ernesto Dal Bo, Ruben Durante, Ed Glaeser, Rick Hornbeck, Kyle Meng, Dina Pomeranz, Sandra Sequeria, Reed Walker, and seminar and conference participants at NBER EEE 2020 Spring Meeting, AEA 2018 Meetings, Univ. of Chicago, Univ of British Columbia, Cornell Univ., Harvard Business School, UC-Berkeley-Haas, INSEAD, PUC-Chile, Sciences Po, Univ. of Hawaii, Univ. of Waikato, UC-Berkeley Econ, UC-Santa Barbara, Vienna Univ. of Economics and Business, Yale University, University of Zurich, and North-east Environmental Economics Workshop, and BREAD conference at Univ. of Maryland for comments. Mobarak acknowledges support from a Carnegie Fellowship Grant ID G-F-17-54329. This study was pre-registered in the AEA RCT Registry: AEARCTR-0000822.

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; Duffo et al., 2013, 2018; Hansman et al., 2018); Actions that affect community health such as open defecation or drunk driving (Banerjee et al., 2017); Or actions that undermine government performance such as corruption or tax evasion (Carrillo et al., 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 collection 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)

¹For 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 et al. (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 et al. (2008), making an initially-effective program completely ineffective in 18 months.

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 behavior², or marketing that appeals to people’s sense of fairness (Hainmueller et al., 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

²For example, Chetty et al. (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 et al. (2015) attempt to change social norms around toilet use (as opposed to directly banning the dangerous practice of open defecation).

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 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

³We use the gendered term “fishermen” because every single fisher we interviewed in Chile was a man.

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 et al., 2010; Lazear, 2006). Also related is the literature on *gaming* incentive schemes where agents adapt to undermine the intent of the regulator (Ederer et al., 2018; Oyer, 1998; Gravelle et al., 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 et al., 2018; Shimshack and Ward, 2005; Gray and Shimshack, 2011; Hansen, 2015; Pomeranz, 2015; Johnson et al., 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 check-points (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

⁴Like 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 et al., 2009; Shimeles et al., 2017; Kollmuss and Agyeman, 2002).

⁵FAO (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

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 et al. (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 2. Section 3 develops the theory of learning. Section 4 describes data collection. Section 5 presents the empirical strategy and results, section 6 documents spillovers and market equilibrium effects, and section 8 concludes the paper.

2 Background and Experimental Design

2.1 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).

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*.

post-harvest sector.”

⁶See <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>

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).

2.2 Supply Chain of Illegal Fish

2.2.1 *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 A.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 in the Appendix describes *caleta* characteristics.

2.2.2 *Ferías*: Outdoor Markets where Hake is Sold

The majority of hake-fish sales to final consumers occur in *ferías*, 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, *ferías* sometimes contain stalls offering fruit and vegetables, clothes and other products.

Ferías 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 *ferías* 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 *ferías*.

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 A.2 and A.3 in the Appendix provide visual examples of ferías and circuits. Table C.1 describes observable characteristics of fish stalls in

ferias.

2.3 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.

2.3.1 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

⁷Santiago 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.

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.

2.3.2 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 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

⁸The 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.

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.

2.3.3 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 A.3 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 2.3.1, 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 A.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 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.

2.3.4 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 1 lists the number of circuits assigned to each of the four treatment cells.

⁹Although 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.

Table 1: Treatment Assignment

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

	High Frequency	Low Frequency	
	Enforcement	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.

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 second panel of the table 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

¹⁰The 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.

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 6.

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.

3.1 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 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, \dots, y_{t-1}) . We assume that the prior $\hat{\theta}_1$ is distributed $\text{Beta}(\alpha_0, \alpha_1)$. Since y_t is Bernoulli, Bayesian updating implies $\hat{\theta}_t$ is distributed $\text{Beta}(\alpha_0 + Y_t, \alpha_1 + t - 1 - Y_t)$, 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 $g : \mathbb{N}_0 \rightarrow (0, 1)$, 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 $\lim_{Y \rightarrow \infty} g(Y) = \bar{g} < 1$.

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:

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

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

- When $g(Y_t) \leq \frac{c}{v}$, the vendor never adopts defensive actions. She sells openly if $\mathbb{E}[\hat{\theta}_t] \leq \frac{v}{\Omega}$, 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 $\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t$; sells hake defensively if $\underline{\delta}_t < \mathbb{E}[\hat{\theta}_t] \leq \bar{\delta}_t$; and does not sell hake if $\mathbb{E}[\hat{\theta}_t] > \bar{\delta}_t$.


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 $\bar{\delta}_t$ 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.

¹¹For 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.

¹²This 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 figure 1(a) 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 1(b) shows the overall sale of hake fish.

Figure 1: Probability of Selling Hake



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

Short versus Long Run In the long-run (as $t \rightarrow \infty$), $\mathbb{E}[\hat{\theta}_t] \rightarrow \theta$, 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 3.2 describes comparative statics of varying enforcement design in this setting.

¹³The specific long-run conditions are discussed in Appendix D.2

In particular, sections 3.2.1 and 3.2.2 discuss the implications of varying the frequency of enforcement visits and the predictability of the enforcement schedule in our field experiment.

3.2 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:

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

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 , a new visit could inadvertently increase the vendor's ability to sell hake illegally. Figure 1(b) 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.

3.2.1 Effects of Frequency of Enforcement Visits


Increasing θ has two effects in equation (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 2 numerically simulates the model for high and low frequency of enforcement over 30 periods, under specific parametric assumptions described in Appendix E. Figures E.2(a) and

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

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

E.2(b) 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.

Figure 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 $\theta^{high} = 0.5, \theta^{low} = 0.3, v \sim U(1/2, 3/2), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7 / (1 + \exp\{-2 \cdot Y + 12\})$, 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.

3.2.2 Predictability of Enforcement Visits

Vendors set up in different *ferias* on different days of the week, as described in section 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 $\theta_t^i (i = 1, 2)$, 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 D.1 details why this corresponds to an optimal belief formation process. We define predictability of the

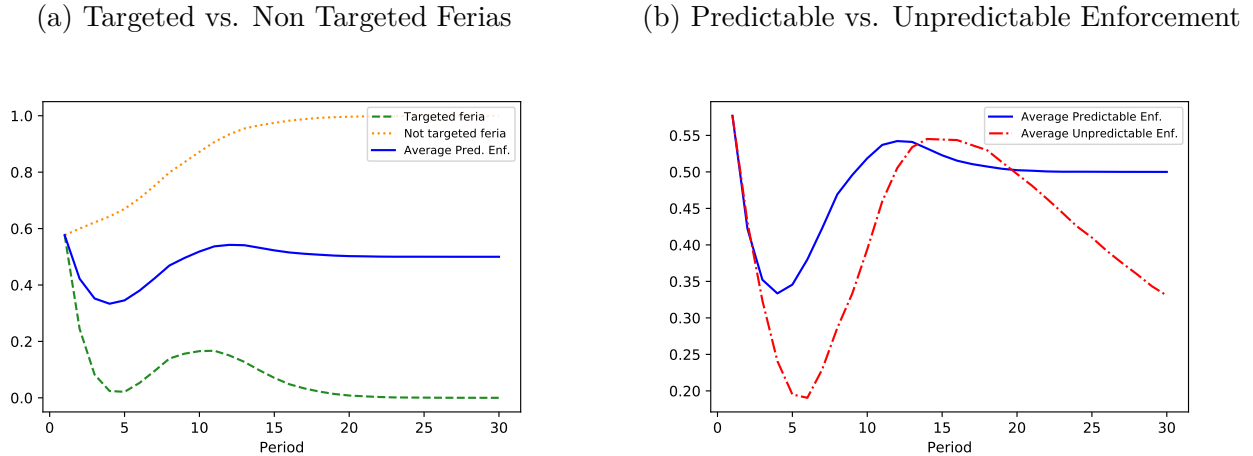
¹⁶Modeling one feria per period (say, f^1 in odd and f^2 in even periods) yields the same qualitative insights.

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

Figure 3: Probability of Selling Hake



Notes: Figure 3(a) 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(b) 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., $\bar{g} = 0.7$.

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(a). Vendor selling strategies diverge between the targeted and non-targeted ferias, and under most functional forms, the unpredictable policy is

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

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.

3.3 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.

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 4 describes the timing of the interventions and data collection activities. In total, seven different data sources are used in the analysis.

Figure 4: Timeline of Interventions and Data Collection



4.1 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.

4.2 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.

4.3 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.

4.4 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 A.4 in Appendix A.2 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 A.8 organizes our interventions and data collection activities along the supply chain for fish.

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 G. 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.2.

5.1 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 :

$$y_{sft} = \beta_0 Post_t + \beta_1 T_c + \beta_2 T_c \times Post_t + \beta_3 y_{sfc0} + X_{ct}' \beta_4 + \varepsilon_{sft} \quad (2)$$

y_{sft} 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_{sft} , 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 :

¹⁸Many 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.

$$y_{ifc} = \gamma_1 T_c + X_{ic}' \Gamma + \epsilon_{ifc} \quad (3)$$

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.

5.2 Treatment Effects on Hake Sales Observed by Mystery Shoppers

Column 1 of Table 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 2 presents the coefficients of interest of regression equation 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

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

¹⁹While 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.

²⁰We 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.

²¹The “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.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.6 shows the results of re-estimating the results in Tables 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

Table 2: Treatment Effects on Hake Sales

VARIABLES	(1)	(2)
	Fresh, Visible Hake	Any Hake Available (Hidden, Frozen, Visible)
Information Campaign Only	-0.133 (0.066)	-0.131 (0.074)
Enforcement Only	-0.178 (0.082)	-0.130 (0.089)
Info Campaign and Enforcement	-0.179 (0.074)	-0.139 (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.

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

on hake sales than the low-saturation treatment.

²²It 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-

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

Table 3: Treatment Effects on Fish Consumption

VARIABLES	(1) Num. Times Hake Purchased	(2) Mention Ban (unprompted)
Information Campaign Only	-0.275 (0.071)	0.146 (0.045)
Enforcement Only	-0.111 (0.049)	0.082 (0.047)
Info Campaign and Enforcement	-0.098 (0.046)	0.107 (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.

Consumer behavior was also indirectly influenced by the enforcement activity. Not only

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.

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.

5.4 Variations in the Design of the Enforcement Strategy

We experimentally manipulated the enforcement schedule in two dimensions: Predictability and Frequency. Table 4 uses the mystery shopper data, and repeats the regression setup of Table 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 3.2.2.

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

Table 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 (0.073)	-0.135 (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 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 et al., 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..


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.

5.5 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.1.4 describes Sernapesca’s enforcement activities in more detail. We organize the data this way to test Prediction (3) from Section 3.2.2 that vendors’ propensity to sell would not remain stable as they learn about enforcement and adapt.

Figure 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 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 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

²³The estimates are obtained from the following regression specification:

$$Y_{sfc} = \sum_{n=0}^N (\beta_n^P \times \mathbf{1}(\#Enf_{ct} = n) \times Pred_c + \beta_n^U \times \mathbf{1}(\#Enf_{ct} = n) \times UnPred_c) + X_{ct}'\Gamma + \varepsilon_{sfc} \quad (4)$$

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

Figure 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.

experimental arm with a predictable visit schedule. This differential effectiveness over time in the predictable arm was also evident in the theoretical simulations [Figures E.4(a) and E.4(b)]. 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.10, vendors who experienced inspections in different ferias on different days of the week (DOWs) reduce hake sales by an extra 9 percentage points ($p\text{-val} < 0.01$) 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.11 and C.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 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, 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 3.2.2. 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 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*

Figure 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.


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 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 6.

Figure 8: Proportion of Defensive Hake



Notes: Figure 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.

Figure B.1 shows that the proportion of hake vendors who adopt these defensive strategies increase week-to-week in response to the enforcement activities. Figure 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 2), and consumer surveys conducted again 6 months after the ban ends (see table C.14) shows that the demand-side effects somewhat persist over the long run.

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 A.4).

6.1 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 5, we re-estimate the effects of predictable and un-predictable enforcement originally reported in Table 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 distance of

²⁴We 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 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.

Table 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 (0.083)	-0.030 (0.069)	-0.076 (0.080)	-0.058 (0.060)
Enforcement on Unpredictable Schedule	-0.157 (0.091)	-0.167 (0.075)	-0.199 (0.084)	-0.177 (0.084)
Spatial Spillover (within 10 km of Treated market)		-0.017 (0.082)		
Social Connection Spill-over (Vendor knows a Treated Vendor)			-0.071 (0.076)	
Supply-Chain Spill-over (Sources from same <i>Caleta</i> as Treated Vendor)				-0.077 (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. .

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

²⁵Using 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.

²⁶Vendors 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).

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.

6.2 Change in Number of Stalls Selling Fish

Our intervention may force some fish vendors to exit the market altogether. Appendix Figure B.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 2. Table C.13 in Appendix C.4 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.

6.3 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 markets 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.

²⁷A 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*.

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 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 (0.105)	0.169 (0.293)	-0.033 (0.147)
Info Campaign	0.043 (0.158)	-0.101 (0.322)	0.343 (0.186)
At least one circuit Enforced and Info Campaign	0.358 (0.128)	0.553 (0.315)	0.173 (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 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

²⁸We could instead define exposure based on the proportion of circuits enforced, and results look similar.

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.

6.4 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.

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.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 7, we study the availability of pomfret (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

The “at least one” formulation is attractive because this indicator evenly divides the sample into equal halves.

²⁹They 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 \(2015\)](#).

Table 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 (0.098)	0.004 (0.035)
Enforcement on Predictable Schedule	0.133 (0.079)	0.027 (0.031)
Enforcement on Unpredictable Schedule	0.115 (0.078)	0.065 (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 \times secret shopper visit, and in the second column is stall \times secret shopper visit \times 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.

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 ($p=0.051$).

³⁰Consumers are more prone to substitute products at similar price levels (see Table C.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.

6.5 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.2). Since pomfret is more often available (and not banned), we instead run regressions to study treatment effects on the price of pomfret.

As a descriptive exercise, Figure B.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 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.

Table 8: Treatment Effect on Fish Prices

VARIABLES	(1) Log Price Pomfret	(2) Log Price Substitute
Information Campaign Only	0.210 (0.109)	0.140 (0.096)
Enforcement Only	-0.017 (0.066)	-0.021 (0.055)
Info Campaign and Enforcement	0.081 (0.065)	0.047 (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* \times *secret shopper visit*, and in the second column is *stall with any substitute available* \times *secret shopper visit* \times *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.

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 6.3 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 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 F.

Table 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 (<i>Overall</i>)	0.13	10,399	\$ 62,900.25	\$ 6.05
<i>Unpredictable</i>	0.192	15,358	\$ 69,190.27	\$ 4.51
<i>Predictable</i>	0.06	4,799	\$ 62,900.25	\$ 13.11
<i>Low Frequency</i>	0.162	12,959	\$ 53,475.84	\$ 4.13
<i>High Frequency</i>	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.1.4, the control group (mistakenly) received a few enforcement visits, the cost is negligible.

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.

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 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 Muñoz-Fraticelli (2013)

³¹An 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.

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

References

- 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.
- 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.
- Christopher Blattman, Donald Green, Daniel Ortega, and Santiago Tobón. Pushing crime around the corner? estimating experimental impacts of large-scale security interventions. 2017.
- 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.
- 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.
- Raj Chetty, Mushfiq Mobarak, and Monica Singhal. Increasing tax compliance through social recognition. *Policy Brief*, 2014.

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

- 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.
- 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 discretion in the enforcement of regulation: Experimental evidence and structural estimates from environmental inspections in india. *Econometrica (forthcoming)*, 2018.
- 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.
- Edward L Glaeser and Andrei Shleifer. The rise of the regulatory state. *Journal of Economic Literature*, 41(2):401–425, 2003.
- 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.
- 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.
- 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.
- 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.

- 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.
- Edward P Lazear. Speeding, terrorism, and teaching to the test. *The Quarterly Journal of Economics*, 121(3):1029–1061, 2006.
- Robert E Lucas. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. Elsevier, 1976.
- 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.
- Dina Pomeranz. No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, 105(8):2539–69, 2015.
- 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.
- 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 Muñoz-Fraticelli. Strategic shortcomings of the dodd-frank act. *The Antitrust Bulletin*, 58(4):617–633, 2013.
- Robert N. Stavins. The problem of the commons: Still unsettled after 100 years. *American Economic Review*, 101(1):81–108, 2011.
- 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.
- WWF. Estimación de la pesca inn en la pesquería de merluza común. *World Wide Fund for Nature*, 2017.

Online Appendix. Not for Publication

A Appendix Figures on the Research Context

A.1 Fishermen Villages

Figure A.1: Fishermen Village (Caleta)



A.2 Outdoor Markets

Figure A.2: Examples of Ferias



Figure A.3: Example of a Circuit





Figure A.3 maps the four ferias that compound one circuit of the city of Los Angeles, VII region.

Figure A.4: Map of Circuits and Caletas



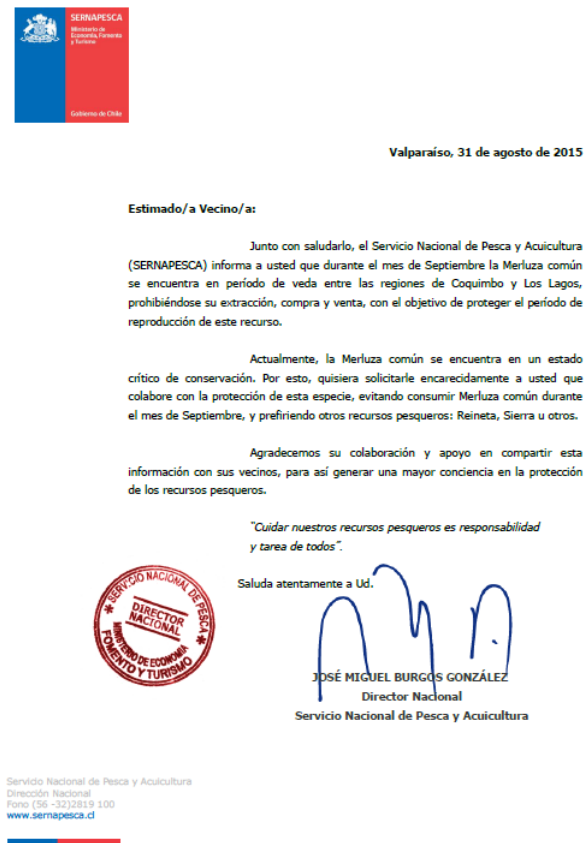
A.3 Interventions

Figure A.5: Flyers




The figure A.5 shows the two types of flyers distributed during the ban period. The message of the one in the left says, “In September respect the Ban”, the one in the right says “This month respect the Ban”.

Figure A.6: Letter to Consumers



The figure A.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 A.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 A.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 A.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 A.8: Interventions and Data Collection at different Points along the Fish Supply Chain



B Appendix Figures on Results

B.1 Adoption of Defensive Actions

Figure B.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 B.1.

Figure B.1: Adoption of Defensive Actions





Figure B.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.

B.2 Number of Stalls

Figure B.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 B.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.

B.3 Prices

Figure B.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 B.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 Appendix Tables

C.1 Descriptive Statistics

C.1.1 Data Collected by Mystery Shoppers

During September 2015, the mystery shoppers interacted with fish-vendors 908 times. The table C.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.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					
<i>No Weight</i>	0.262	0.440	0	1	848
<i>Mechanical Weight</i>	0.410	0.492	0	1	848
<i>Digital Weight</i>	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.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.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).

C.1.2 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 describes the main variables collected in the survey.

Table C.3: Fishermen Characteristics

Variable	Mean	SD	Min	Max	N
<i>Fisherman Boat characteristics:</i>					
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
<i>Union Participation:</i>					
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
<i>Number of Days that Goes Fishing Every Week:</i>					
Summer	5.01	1.47	1	7	226
Winter	2.25	1.14	0	7	227
<i>Number of Boats in the Caleta:</i>					
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
<i>Top 3 Most Captured Fishes in the Caleta:</i>					
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
<i>Usual Buyer of the Fish at the Dock:</i>					
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.

C.1.3 Consumer Mobility Between Neighborhoods

The table C.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.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.

C.1.4 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.³³ In total, 230 visits were carried out, equivalent to 659 stall-inspections in 62 circuits. Based on the inspectors' reports,

³³These 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.

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.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.

Table C.5: Implementation of the Interventions

Treatment Assignment	(1)	(2)	(3)	(4)	(5)	(6)
	Number of Visits	Number of Different Days of the Week Visited	Circuit Visited at Least Once	At Least One Visit		N
				Number of Visits	Number of Different Days of the Week Visited	
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.

C.2 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.6 shows balance tests across the main treatment arms. Tables C.7 and C.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.

C.2.1 Balance Tables

The table C.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.7 and C.8 present the same estimates but decomposing by the enforcement variations: predictability and frequency.

Table C.6: Randomization Balance on Observables

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Control	Info Campaign		Enforcement		Enforc. and Info Camp.	
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573 (0.497)	0.644 (0.484)	0.083 [0.740]	0.489 (0.501)	-0.069 [0.574]	0.509 (0.501)	0.013 [0.920]
Distance to Closest Caleta (kms)	16.507 (25.082)	11.572 (9.503)	3.465 [0.516]	14.863 (22.626)	-3.606 [0.386]	26.425 (29.037)	2.245 [0.682]
Poverty Rate Municipality	19.006 (4.780)	17.567 (3.313)	-2.148 [0.244]	18.026 (5.483)	-1.079 [0.412]	16.734 (7.549)	-0.316 [0.851]
Av. Monthly Income Municipality (USD)	791.767 (149.808)	875.475 (172.858)	18.446 [0.846]	790.514 (140.334)	-2.506 [0.953]	830.683 (139.251)	20.464 [0.673]
Delinquency Rate Municipality	0.038 (0.015)	0.029 (0.002)	-0.013 [0.016]	0.036 (0.015)	-0.001 [0.835]	0.034 (0.009)	-0.004 [0.480]
Rain Indicator	0.290 (0.456)	0.133 (0.344)	-0.124 [0.455]	0.178 (0.383)	-0.114 [0.318]	0.142 (0.349)	-0.122 [0.301]
Average Temperature (Celsius)	12.200 (2.281)	12.126 (2.087)	0.081 [0.942]	11.993 (2.021)	-0.192 [0.797]	11.936 (2.196)	-0.688 [0.346]
Joint Significance							
<i>F statistic</i>			0.609		1.094		0.816
<i>p-value</i>			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.7: Randomization Balance: Enforcement Predictability

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Control	Info Campaign		Enforc: Predictable Schedule		Enforc: Unpredictable Schedule	
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573 (0.497)	0.644 (0.484)	0.080 [0.749]	0.396 (0.490)	-0.148 [0.270]	0.575 (0.495)	0.045 [0.710]
Distance to Closest Caleta (kms)	16.507 (25.082)	11.572 (9.503)	3.822 [0.479]	17.286 (23.417)	1.704 [0.678]	20.487 (27.384)	-4.926 [0.333]
Poverty Rate Municipality	19.006 (4.780)	17.567 (3.313)	-2.173 [0.240]	17.130 (6.185)	-1.739 [0.212]	17.890 (6.450)	-0.075 [0.960]
Av. Monthly Income Municipality (USD)	791.767 (149.808)	875.475 (172.858)	18.828 [0.843]	809.457 (153.747)	-1.068 [0.981]	801.805 (130.506)	9.235 [0.834]
Delinquency Rate Municipality	0.038 (0.015)	0.029 (0.002)	-0.013 [0.015]	0.036 (0.015)	-0.002 [0.648]	0.035 (0.012)	-0.001 [0.782]
Rain Indicator	0.290 (0.456)	0.133 (0.344)	-0.126 [0.446]	0.117 (0.322)	-0.158 [0.162]	0.202 (0.402)	-0.080 [0.488]
Average Temperature (Celsius)	12.200 (2.281)	12.126 (2.087)	0.076 [0.946]	12.058 (2.057)	-0.170 [0.824]	11.904 (2.107)	-0.491 [0.498]
Joint Significance							
<i>F statistic</i>			0.609		1.954		1.717
<i>p-value</i>			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.8: Randomization Balance: Enforcement Frequency

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	Control	Info Campaign		Enforc: High Frequency		Enforc: Low Frequency	
	Mean	Mean	Diff	Mean	Diff	Mean	Diff
Indicator Fixed Stalls	0.573 (0.497)	0.644 (0.484)	0.098 [0.696]	0.372 (0.484)	-0.164 [0.211]	0.591 (0.492)	0.039 [0.753]
Distance to Closest Caleta (kms)	16.507 (25.082)	11.572 (9.503)	4.009 [0.451]	16.224 (23.309)	-5.302 [0.330]	21.269 (27.292)	0.600 [0.885]
Poverty Rate Municipality	19.006 (4.780)	17.567 (3.313)	-1.975 [0.279]	16.242 (7.011)	-2.256 [0.101]	18.564 (5.576)	0.140 [0.922]
Av. Monthly Income Municipality (USD)	791.767 (149.808)	875.475 (172.858)	17.128 [0.857]	821.315 (142.660)	23.237 [0.587]	792.779 (138.930)	-8.734 [0.840]
Delinquency Rate Municipality	0.038 (0.015)	0.029 (0.002)	-0.013 [0.014]	0.036 (0.013)	-0.000 [0.973]	0.035 (0.014)	-0.003 [0.547]
Rain Indicator	0.290 (0.456)	0.133 (0.344)	-0.128 [0.439]	0.193 (0.396)	-0.077 [0.491]	0.143 (0.350)	-0.143 [0.215]
Average Temperature (Celsius)	12.200 (2.281)	12.126 (2.087)	0.050 [0.964]	12.127 (2.204)	-0.192 [0.799]	11.853 (1.984)	-0.447 [0.545]
Joint Significance							
<i>F statistic</i>			0.609		2.016		2.090
<i>p-value</i>			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.9 shows the coefficients of regression 2. The first three rows show the differences in the sale of illegal hake in the pre-intervention period. The interaction of “ \times 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.9: Treatment Effects on Hake Sales

VARIABLES	(1)	(2)
	Fresh, Visible Hake	Any Hake Available (Hidden, Frozen, Visible)
Information Campaign Only	0.080 (0.056)	0.029 (0.058)
Enforcement Only	0.114 (0.070)	0.092 (0.060)
Information Campaign and Enforcement	0.078 (0.070)	0.100 (0.065)
Information Campaign Only \times Post	-0.133 (0.066)	-0.131 (0.074)
Enforcement Only \times Post	-0.178 (0.082)	-0.130 (0.089)
Info Campaign and Enforcement \times Post	-0.179 (0.074)	-0.139 (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.

C.3 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.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.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 \times 2nd half” shows the same vendor’s decision to sell on a weekday in which he did not experience

Table C.10: Hake Available based on the Number of different ferias and Days of the Week visited

VARIABLES	(1)	(2)	(3)
	Any Hake Available		
N Ferias Visited	0.041 (0.030)		0.039 (0.033)
N Ferias Visited \times Second Half	-0.091 (0.023)		-0.081 (0.035)
N DOWs Visited		0.030 (0.055)	0.014 (0.056)
N DOWs Visited \times Second Half		-0.098 (0.073)	-0.037 (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.

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.

We consider this evidence as only suggestive and placed it in the appendix, because Ser-napesca 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.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 “... \times 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

Table C.11: Hake Sale across DOWs and Ferias within Circuit

VARIABLES	(1)	(2)	(3)	(4)
	Any Hake Available			
	Full Sample		Circuits that Rotate between More Than 2 Ferias	
Feria Not Visited	-0.209 (0.157)		-0.299 (0.182)	
Feria Not Visited \times Second Half	0.058 (0.086)		0.267 (0.130)	
DOW Not Visited		-0.252 (0.131)		-0.303 (0.161)
DOW Not Visited \times Second Half		0.192 (0.104)		0.286 (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.

Table C.12: Number of Stalls Hake Sale across DOWs and Ferias within Circuit

VARIABLES	(1)	(2)	(3)	(4)
	Number of Stalls			
	Full Sample		Circuits that Rotate between More Than 2 Ferias	
Feria Not Visited	-0.510 (0.320)		-0.956 (0.368)	
Feria Not Visited \times Second Half	0.354 (0.464)		1.331 (0.771)	
DOW Not Visited		-0.024 (0.220)		-0.241 (0.312)
DOW Not Visited \times Second Half		0.317 (0.330)		0.995 (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.

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.

C.4 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.³⁴ 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.^{35 36}

³⁴We allow the number of missing stalls to be non-integer, and negative if the number of stalls increased.

³⁵If 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.

³⁶Since 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.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 (0.060)	-0.115 (0.065)
Enforcement Only	-0.190 (0.082)	-0.141 (0.091)
Info Campaign and Enforcement	-0.156 (0.084)	-0.130 (0.104)
Panel B: Variation in Predictability of Enforcement		
Info Campaign Only	-0.111 (0.062)	-0.121 (0.064)
Enforcement on Predictable Schedule	-0.091 (0.073)	-0.061 (0.087)
Enforcement on Unpredictable Schedule	-0.246 (0.089)	-0.197 (0.100)
Panel C: Variation in Frequency of Enforcement		
Info Campaign Only	-0.113 (0.062)	-0.121 (0.064)
High Frequency Enforcement	-0.086 (0.092)	-0.092 (0.101)
Low Frequency Enforcement	-0.184 (0.086)	-0.148 (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 comparing the number of stalls per circuit before and after the interventions). The panel A describes the same specification presented in Table 2. Panels B and C follow the same specification than Table 4. Probit regression marginal effects are reported. Robust standard errors clustered by circuit in parentheses.

C.5 Treatment Effects Six Months After the Ban Period

Table C.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.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 (0.106)	-0.419 (0.234)	-0.121 (0.100)
Enforcement Only	0.018 (0.078)	-0.106 (0.205)	0.093 (0.069)
Info Campaign and Enforcement	-0.017 (0.081)	-0.197 (0.246)	0.017 (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.

C.6 Alternative Definition Information Campaign Treatment

Tables C.15 and C.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.15: Treatment Effect on Hake Availability

VARIABLES	(1)	(2)
	Fresh, Visible Hake	Any Hake Available (Fresh-Visible, Hidden or Frozen)
Information Campaign Only	-0.082 (0.064)	-0.070 (0.071)
Enforcement Only	-0.157 (0.079)	-0.101 (0.094)
Info Campaign and Enforcement	-0.169 (0.079)	-0.121 (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 the interactions $T_c \times Post_t$ are reported. Robust standard errors are clustered by circuit and presented in parentheses.

Table C.16: Treatment Effect on Hake Sales by Enforcement Strategy

VARIABLES	(1)	(2)
	Any Hake	Available
Info Campaign Only	-0.073 (0.071)	-0.073 (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 the interactions $T_c \times Post_t$ are reported. Robust standard errors are clustered by circuit and presented in parentheses.

D Appendix: Theoretical Model

D.1 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 assume is the underlying distribution determining the probability of a visit in each feria.³⁷ 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 $\hat{p}_0 \sim \text{Dirichlet}((\beta_i)_{i=1}^4)$. Finally, we denote by $\theta = (\theta^1, \theta^2)$ the real probability of visits, which we call the *visit policy*.³⁸ 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*

$$\begin{aligned}\hat{\theta}_t^1 &\sim \text{Beta}\left(\alpha_2 + \alpha_4 + Y_t^1; \alpha_1 + \alpha_3 + t - 1 - Y_t^1\right) \\ \hat{\theta}_t^2 &\sim \text{Beta}\left(\alpha_3 + \alpha_4 + Y_t^2; \alpha_1 + \alpha_2 + t - 1 - Y_t^2\right)\end{aligned}$$

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.

D.2 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

$$\mathbb{E}[\hat{\theta}] \leq \frac{1}{\Omega} \min \left\{ v; \frac{c}{g} \right\}.$$

³⁷Note 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.

³⁸As the distribution of z_t is stationary, the probabilities of visits in both ferias also are.

- 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 $\bar{\delta} = \frac{v-c}{\Omega(1-g)}$)

$$\underline{\delta} < \mathbb{E}[\hat{\theta}] \leq \bar{\delta}.$$

- 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

$$\mathbb{E}[\hat{\theta}] > \frac{1}{\Omega} \max \left\{ v; \frac{v-c}{1-g} \right\}.$$

First, note that the conditions $\underline{\delta} < \bar{\delta}$, $v > \frac{c}{g}$, 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 $\underline{\delta} \geq \bar{\delta}$ 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 from the comparison of the three options.

[Long-run Comparative Statics] Define $\underline{\delta}_\infty = \frac{c}{\Omega \bar{g}}$ and $\bar{\delta}_\infty = \frac{v-c}{\Omega(1-\bar{g})}$. First, note that Assumption $\bar{g} > c/v$ implies that $\underline{\delta}_\infty < \bar{\delta}_\infty$. As (a.s.) $\underline{\delta}_t \rightarrow \underline{\delta}_\infty$ and $\bar{\delta}_t \rightarrow \bar{\delta}_\infty$, 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

$$\theta \leq \bar{\delta}_\infty = \frac{v-c}{\Omega(1-\bar{g})}.$$

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 $Z_t = (Z_t^j)_{j=1}^4$, Bayesian updating implies that

$$\hat{p}_t \sim \text{Dir}(\alpha + Z_t)$$

As $Z_t^2 + Z_t^4 = Y_t^1$, the previous distribution implies that the vector

$$(\hat{p}_t^1, \hat{p}_t^2 + \hat{p}_t^4, \hat{p}_t^3) = (\hat{p}_t^1, \hat{\theta}_t^1, \hat{p}_t^3) \sim \text{Dir}(\alpha^1 + Z_t^1, \alpha^2 + \alpha^4 + Y_t^1, \alpha^3 + x_t^3).$$

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

$$\hat{\theta}_t^1 \sim \text{Beta}(\alpha^2 + \alpha^4 + Y_t^1, \alpha^1 + \alpha^3 + t - 1 - Y_t^1).$$

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 $\bar{\delta}_t$. Hence, in the long-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 $\bar{\delta}_t$ in both ferias, the inspector should choose a targeted policy to prevent illegal sales in at least one feria

[Proof of Corollary 1] We analyze the different cases characterized in Proposition 1. For the analysis we use that $\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t \iff \Omega\mathbb{E}[\hat{\theta}_t] \leq c/g(Y_t)$, and $\mathbb{E}[\hat{\theta}_t] \leq \bar{\delta}_t \iff \Omega\mathbb{E}[\hat{\theta}_t](1-g(Y_t))+c \leq v$.

- The agent does not sell in two cases

1. If $v \leq c/g(Y_t)$, the vendor does not sell if $v < \Omega\mathbb{E}[\hat{\theta}_t]$. This happens with probability

$$\begin{aligned} \mathbb{P}\left(v \leq \frac{c}{g(Y_t)}\right) \mathbb{P}\left(v < \Omega\mathbb{E}[\hat{\theta}_t] \middle| v \leq \frac{c}{g(Y_t)}\right) &= \mathbb{P}\left(\left(v < \Omega\mathbb{E}[\hat{\theta}_t]\right) \left(v \leq \frac{c}{g(Y_t)}\right)\right) \\ &= \mathbb{P}\left(v < \min\left\{\frac{c}{g(Y_t)}; \Omega\mathbb{E}[\hat{\theta}_t]\right\}\right) \\ &= F\left(\min\left\{\frac{c}{g(Y_t)}; \Omega\mathbb{E}[\hat{\theta}_t]\right\}\right). \end{aligned}$$

2. If $v > c/g(Y_t)$ the vendor does not sell if $\mathbb{E}[\hat{\theta}_t] > \bar{\delta}_t$.³⁹ The probability of this is

$$\begin{aligned} \mathbb{P}\left(v > \frac{c}{g(Y_t)}\right) \mathbb{P}\left(v < \Omega\mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c \mid v > \frac{c}{g(Y_t)}\right) &= \\ &= \mathbb{P}\left(\left(v < \Omega\mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c\right) \mid \left(v > \frac{c}{g(Y_t)}\right)\right) \\ &= \mathbb{P}\left(\frac{c}{g(Y_t)} < v < \Omega\mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c\right) \\ &= \max\left\{F\left(\Omega\mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c\right) - F\left(\frac{c}{g(Y_t)}\right); 0\right\}. \end{aligned}$$

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 $v \leq c/g(Y_t)$, the vendor does not sell if $v > \Omega\mathbb{E}[\hat{\theta}_t]$. This happens with probability

$$\begin{aligned} \mathbb{P}\left(v \leq \frac{c}{g(Y_t)}\right) \mathbb{P}\left(v > \Omega\mathbb{E}[\hat{\theta}_t] \mid v \leq \frac{c}{g(Y_t)}\right) &= \mathbb{P}\left(\left(v > \Omega\mathbb{E}[\hat{\theta}_t]\right) \mid \left(v \leq \frac{c}{g(Y_t)}\right)\right) \\ &= \mathbb{P}\left(\Omega\mathbb{E}[\hat{\theta}_t] < v < \frac{c}{g(Y_t)}\right) \\ &= \max\left\{F\left(\frac{c}{g(Y_t)}\right) - F\left(\Omega\mathbb{E}[\hat{\theta}_t]\right); 0\right\}. \end{aligned}$$

2. If $v > c/g(Y_t)$ the vendor sells openly if $\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t$.⁴⁰ The probability of this is

$$\begin{aligned} \mathbb{P}\left(v > \frac{c}{g(Y_t)}\right) \mathbb{P}\left(\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t \mid v > \frac{c}{g(Y_t)}\right) &= \mathbb{P}\left(\left(\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t\right) \mid \left(v > \frac{c}{g(Y_t)}\right)\right) \\ &= \mathbb{P}\left(\frac{c}{g(Y_t)} < v\right) \mathbb{1}_{\{\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t\}} \\ &= \left(1 - F\left(\frac{c}{g(Y_t)}\right)\right) \mathbb{1}_{\{\Omega\mathbb{E}[\hat{\theta}_t] \leq \frac{c}{g(Y_t)}\}}. \end{aligned}$$

The share of vendors who sell openly α_{SO} is the sum of these two probabilities.

- The vendor sells defensively only if $\bar{\delta}_t < \mathbb{E}[\hat{\theta}_t] \leq \bar{\delta}_t$.⁴¹ The share of vendors who sell

³⁹If $v > c/g$ and $\mathbb{E}[\hat{\theta}_t] > \bar{\delta}_t$, the condition $\mathbb{E}[\hat{\theta}_t] > \underline{\delta}_t$ is necessarily satisfied.

⁴⁰If $v > c/g$ and $\mathbb{E}[\hat{\theta}_t] \leq \underline{\delta}_t$ the condition $\mathbb{E}[\hat{\theta}_t] \leq \bar{\delta}_t$ is necessarily satisfied.

⁴¹Recall that the conditions $\underline{\delta}_t < \bar{\delta}_t$ and $v > c/g(Y_t)$ 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 $\Omega\mathbb{E}[\hat{\theta}_t] \leq c/g(Y_t) \iff \Omega\mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) - c \leq c/g(Y_t)$ and replace the corresponding values in each case.

E 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 E.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 E.1 is that vendors' behavior varies over time; the same policy evaluated in different moments may yield different results.

Figure E.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 $\theta = 0.4, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7 / (1 + e^{-8 \times Y + 28})$, 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 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 E.2(a) and E.2(b) 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 E.2: Probability of Selling Hake

(a) Low Frequency Enforcement



(b) High Frequency Enforcement




Figure E.2(a) and E.2(b) describe vendors' decision on whether and how to sell. This simulation uses the same parameters than previous graph: $\theta^{high} = 0.5, \theta^{low} = 0.3, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7 / (1 + e^{-2 \times Y + 12})$, 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 E.3(a) and E.3(b) 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 E.3: Probability of Selling Hake





Figure E.3(a) and E.3(b) describe vendors' decision vary depending on the feria they are selling. The model assumes vendors alternate between targeted 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 $\theta^{high} = 0.4$, $\theta^{low} = 0.25$, $v \sim U(0.5, 1.5)$, $c = 0.1$, $\Omega = 18$, $\theta_1 = 0.05$, $g(Y) = 0.7 / (1 + e^{-8 \times Y + 28})$, i.e., $\bar{g} = 0.7$.

The figures E.4(a) and E.4(b) compare the average probability of selling hake under predictable and unpredictable enforcement using the same enforcement capacity. As discussed in section 3.2.2, 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 E.4: Probability of Selling Hake

(a) Low Frequency Enforcement



(b) High Frequency Enforcement




Figure E.4(a) and E.4(b) describe vendors' decision vary depending on the feria they are selling. The model assumes vendors alternate between targeted 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 $\theta^{high} = 0.4, \theta^{low} = 0.25, v \sim U(0.5, 1.5), c = 0.1, \Omega = 18, \theta_1 = 0.05, g(Y) = 0.7 / (1 + e^{-8 \times Y + 28})$, 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 $[\underline{v}, \bar{v}]$. Assume $c < \underline{v} \leq \bar{v} < \Omega$. The applying 1 we get the following result

Corollary 1 *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*

$$\alpha_{NS} = \begin{cases} F(\Omega \mathbb{E}[\hat{\theta}_t]) & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] \leq c/g(Y_t) \\ F(\Omega \mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c) & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] > c/g(Y_t); \end{cases}$$

$$\alpha_{SO} = \begin{cases} 1 - F(\Omega \mathbb{E}[\hat{\theta}_t]) & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] \leq c/g(Y_t) \\ 0 & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] > c/g(Y_t); \text{ and} \end{cases}$$

$$\alpha_{SD} = \begin{cases} 0 & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] \leq c/g(Y_t) \\ 1 - F(\Omega \mathbb{E}[\hat{\theta}_t](1 - g(Y_t)) + c) & \text{if } \Omega \mathbb{E}[\hat{\theta}_t] > c/g(Y_t). \end{cases}$$

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, $\lim_{x \rightarrow \infty} g(x) = \bar{g}$, and the parameters a and b dictate the speed of convergence of the function. Other functional forms yield the same qualitative results. The figures E.5(a) and E.5(b) 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 E.5: Beliefs Updating and Learning Curve




Figure E.5(a) and E.5(b) how vendors learn about loopholes and update the probability of a visit. These figures use the same inputs than other simulations, i.e., $\bar{g} = 0.7, a = 8, b = 28, \theta_1 = 0.05 (\alpha_0 + \alpha_1 = 40)$

F Further Details on the Cost-Effectiveness Analysis

The section 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 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.

G 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.