Essays on Public Economics

by

Cristobal Otero

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

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Emmanuel Saez, Co-chair
Professor Gabriel Zucman, Co-chair
Professor Sydnee Caldwell
Professor Frederico Finan

Spring 2023

Essays on Public Economics

Abstract

Essays on Public Economics

by

Cristobal Otero

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Emmanuel Saez, Co-chair

Professor Gabriel Zucman, Co-chair

This dissertation explores empirically the role of government in addressing contemporary public health pressing issues. The first chapter studies the entrance of state-owned pharmacies to provide access to affordable pharmaceutical drugs. This is joint work with Juan Pablo Atal from University of Pennsylvania, José Ignacio Cuesta from Stanford University, and Felipe González from Queen Mary University of London. We find that public pharmacies sell the same drugs at a third of private pharmacy prices, because of stronger upstream bargaining and downstream market power in the private sector but are of lower quality. Leveraging the decentralized entry of public pharmacies to local markets in Chile, we show that public pharmacies induced market segmentation and price increases in the private sector, benefiting the switchers to the public option but harming the stayers. We conclude that the countrywide entry of public pharmacies would reduce yearly consumer drug expenditure and significantly outweigh the costs of the policy.

In the second chapter, in joint work with Pablo Muñoz from Universidad de Chile, we ask how governments can improve healthcare provision in public hospitals. To this end, we study a reform in Chile that aimed to improve public service provision by dramatically changing the way government institutions recruit high-ranking civil servants. We focus on the impact of the policy on public hospital performance and examine the underlying mechanisms through which public managers affect public health outcomes. The paper shows that the policy reduced hospital mortality around 8%, an effect that persisted after three years. We also find that the policy changed the pool of CEOs by displacing older doctors with no management training in favor of younger CEOs with either undergraduate degrees in management or doctors with master's degrees or diplomas in management. We find that the reform affected hospital mortality mostly when newly appointed managers had management studies, who introduced more efficient use of medical resources and better personnel practices.

The third chapter studies an innovative nationwide policy that mandates the use of warning labels on products whose sugar or calorie concentration exceeds certain thresholds to address increasing obesity. The policy was first passed in Chile and has been widely repli-

cated in several countries such as Argentina, Brazil, Canada, Mexico, Israel, among others. In a joint project with Nano Barahona and Sebastián Otero, both from UC Berkeley, we partnered with Walmart-Chile to study the effects of the policy Chile. We find that consumers substituted from labeled to unlabeled products, a pattern mostly driven by products that consumers mistakenly believe to be healthy. On the supply side, we document substantial reformulation of products and bunching at the thresholds. Next we develop and estimate an equilibrium model of demand for food and firms' pricing and nutritional choices. The main finding is that that food labels increase consumer welfare, an effect that is enhanced by firms' responses. We conclude that under optimal policy thresholds, food labels and sugar taxes generate similar gains in consumer welfare, but food labels benefit the poor relatively more.

For Josefina, Sam & Martin.

Contents

List of Figures

List of Tables

Acknowledgments

Coauthor information: Chapter 1 is coauthored with Juan Pablo Atal from University of Pennsylvania, José Ignacio Cuesta from Stanford University, and Felipe González from Queen Mary University of London. Chapter 2 is coauthored with Pablo Muñoz from Universidad de Chile and former graduate student in economics at UC Berkeley. Chapter 3 is coauthored with Nano Barahona, and Sebastián Otero. Nano is a current assistant professor at UC Berkeley and Sebastián is a postdoctoral researcher at UC Berkeley.

Acknowledgements: Josefina and our two kids, Sam and Martin, have been my main support and inspiration during the last six years at Berkeley. While pursuing a PhD can be an emotional rollercoaster, my family has been a constant reminder of what truly matters. They teach me every day that the journey might be even more important than the destination, and as long as we're traveling together, that's all I need. In hindsight, I'm even grateful for the late nights and sleepless hours that Sam and Martin so skillfully provided. I don't want to suggest that having kids during the PhD makes it easier, but my boys kept me grounded and gave me perspective and discipline. And they make me happy every single day. I also have to acknowledge that part of my dissertation is literally Josefina's work. We worked hand by hand collecting data and even organizing field work. She is my greatest fan and the best journey companion I could've ever wished for. I'm looking forward to our next adventure and to keep learning and having fun together. I love you.

At Berkeley, I was fortunate to find a community of mentors and friends. I initially chose the university to connect with Emmanuel Saez and Gabriel Zucman, whose work I've long admired. The way their academic research informs and feeds policy and their belief that institutions are not immutable forces of nature but can be improved through political decisions, have been constant inspirations. I arrived with the simple expectation of taking classes with them, never imagining both would co-chair my dissertation committee and provide continuous guidance. Both are not only exceptionally talented, but also down-to-earth, thoughtful, and caring. Emmanuel has a unique ability to express complex ideas in a simple and intuitive manner, and he taught me the value of clarity and transparency in research. I'm grateful for his support that went beyond the strict academic boundaries. In Gabriel, I found a mentor and a friend. He consistently supported and encouraged me throughout the entire PhD and he was key for my personal and academic success during these years. I found him to be an extremely generous and kind person. His open-door policy extended even to my family into his home. I will never forget the unexpected ways in which he supported me when my second child was born, and I'll always be grateful. Gabriel also gave me the opportunity to work closely with him, and consistently showed that he trusted in me and valued my work. What is really surprising is that my experience is not an exception. Gabriel and Emmanuel have created a network of students and scholars who have flourished by their side. I'm extremely thankful and grateful, and I really hope that our paths will keep crossing in the future.

Sydnee Caldwell also played an essential role in this journey. Prior to her arrival at Berkeley, I knew she would be an ideal addition to my orals committee. Since her official appointment began a few months after my orals, I had to seek special permission from the Graduate Division to include her on my committee. As a result, I had the distinction of being the first student for whom she participated in an orals committee at Berkeley. And I was really lucky. Sydnee took the job as if she was referring to a paper for a top journal and sent me detailed and sharp comments and feedback. She continued to support, encourage, and advise me for the following years, and I'm certain that she is the main reason why my job market paper made it from the orals to my main dissertation (Chapter 2).

Fred Finan was also important, especially because he always believed in me, and that matters. Fred was on the PhD selection committee when I applied to the PhD and gave me a chance to study at Berkeley despite my non-traditional path to an economics PhD, which included coming straight from a master's in philosophy. During the PhD, he generously offered advice, sat on my orals committee, and was kind enough to write me a recommendation letter. In his way, he was always present in my academic adventure, and his willingness to vouch for me gave me much-needed reassurance in this type of endeavor.

I would also like to send a big thank you to my parents. I am really lucky to have them. I was raised in an environment full of love and support that allowed me to develop my talents and gave me the means to achieve my goals. My mom taught me discipline, opened me to the world through books and reading, and constantly nurtured my curiosity. My dad taught me that squares can be circled and that your imagination is the limit. During my PhD, my dad visited regularly and spent a lot of time simply being part of Sam and Martin's life. Every visit was a source of joy for us and especially for Sam. He also took care of every detail that I usually overlooked, but that would make life easier. Pursuing a PhD becomes easier when you receive such constant and significant support.

My brother, Sebastián, was exceptionally important. Part of the reason I pursued a PhD was seeing him as a role model. I'll always remember our trip to LA before starting the PhD and will always be grateful for his continuous and generous support throughout my PhD, especially during the job market process. I'm sure I wouldn't have landed a job without Seba's numerous pep talks and advice during the most challenging moments. Our work (some of which we coauthored) and a lot of luck put us together in New York, and I'm excited to navigate this new adventure side by side again. I'm also grateful to Anna Popova for her patience and constant support in reviewing my work. Anna taught me a valuable lesson: our eyes process writing before our brain comprehends it.

Life is better when shared, and this is also true in academia. My coauthors were central to my PhD and professional development, and I've been extremely fortunate to learn from them. Pablo Muñoz is bright, hardworking, and provides the perfect amount of judgment needed in long and grinding academic projects. Nano Barahona showed me what constant and thorough work looks like and that brain without muscle is unlikely to succeed. Felipe González before being my coauthor, was my TA back in Chile. I think part of my academic curiosity is due to his passion for research and teaching. He is a wonderful scholar and friend. José Ignacio Cuesta is the kind of person who is always spot on and makes everything look straightforward. It's like having a coauthor who is simultaneously your referee. No wonder why he is doing so well. Juan Pablo Atal taught me how to use math in a creative and simple way to convey clear economic intuitions. He was also the first Chilean I met who was working in the US academia, which was an important inspiration during my PhD.

My office mates in room 610 were another source of support and joy. Damián Vergara, Patrick Kennedy, Harry Wheeler, and Felipe Arteaga played an important role in their own way. I'll keep a special place in my heart for the time I spent sharing the office with them (even if their presence was spiritual in some cases). My friends in Chile also have supported me along the way and they keep inspiring me to challenge my certainties and to work for a greater purpose. I cannot finish without mentioning Ponchito Mora, who never missed an up or down; he walked with me in the tough moments and celebrated me in the good ones.

I would finally like to thank the financial support from several donors and institutions that made my work possible. I'm thankful to Conicyt Becas Chile, the Institute for Research on Labor and Employment, the Center for Equitable Growth at UC Berkeley, the UC Berkeley James M. and Cathleen D. Stone Center, and the Robert D. Burch Center for Tax Policy & Public Finance.

I came to Berkeley with the goals of learning, growing, and connecting, and I'm grateful for the many opportunities and relationships that have made these past six years genuinely unforgettable.

Chapter 1

The Economics of the Public Option: Evidence from Local Pharmaceutical Markets

1.1 Introduction

State-owned firms compete with the private sector in education, healthcare, insurance, and basic services, among others. Supporters of the public option argue that it helps discipline markets that fail to provide enough incentives for private competition, because of either information asymmetries, market power, collusive behavior, or other market failures (Atkinson and Stiglitz, 1980). In contrast, critics argue that state-owned firms might be inefficient, provide low quality, or be captured by political interests (Shleifer and Vishny, 1994; Shleifer, 1998). Estimating the equilibrium effects of the public option has been difficult due to the lack of exogenous variation in the extent of public competition and the scarcity of contexts that allow evaluation of its distributional and political consequences.

In this paper, we study the decentralized and large-scale entry of public retail pharmacies in Chile, where pharmacies managed by local governments entered 146 of the 344 counties between 2015 and 2018. Public pharmacies emerged as nonprofit competition to a fully deregulated and highly concentrated private retail market characterized by high prices.[1] Public pharmacies sell drugs at prices that are 34 percent of those charged by their private counterparts. These low prices are possible because private pharmacies hold substantial market power and public pharmacies have a cost advantage. However, public pharmacies are of lower quality than their private counterparts: They require consumers to travel more than two times more, carry less product variety, and have more restrictive operating hours and longer waiting times.

To estimate the impacts of public pharmacies, we combine quasi-experimental approaches

[1] Chile has relatively high drug prices and high out-of-pocket spending as a share of health expenditures compared with other OECD countries (OECD, 2015).

with a field experiment to study market outcomes and political preferences. The quasi-experiment exploits the staggered entry of public pharmacies across counties. To support this design, we show that the timing of entry was unrelated to baseline differences or pre-trends in local market attributes. Moreover, anecdotal evidence suggests that the timing of entry of public pharmacies depended partly on unexpected delays in the bureaucratic procedure for obtaining sanitary permits. The field experiment consisted of an informational intervention with consumers, which we conducted during the weeks preceding the 2016 local election in counties with public pharmacies. The treatment covered the existence, location, low prices, and low convenience of public pharmacies. We surveyed consumers before the intervention and two months after, collecting data about drug shopping behavior and political participation.

We begin by estimating how the entry of public pharmacies impacted private-sector market outcomes. We exploit the staggered entry of public pharmacies and drug-level data to estimate their impact on private pharmacy prices and sales. Eighteen months after opening, the average public pharmacy had shifted 4 percent of sales away from private pharmacies. The decrease in sales was concentrated among drugs that target chronic conditions. We also find a *positive* and growing effect of public pharmacies on private sector prices: By the end of our sample period, the entry of public pharmacies had induced private pharmacies to increase their prices by 1 percent. We interpret this positive price effect as evidence that this low-price and low-quality public option generated market segmentation. In particular, private pharmacies responded to a shift of relatively price-sensitive consumers toward public pharmacies—and thus a less elastic residual demand—by increasing prices. This result is consistent with theoretical research on the potential for price-increasing competition (Chen and Riordan, 2008). A simple model of competition with differentiated firms rationalizes the lack of a stronger demand shift to public pharmacies, despite their low relative prices, as a result of low relative quality. These results show that public pharmacies generated winners and losers as a consequence of their equilibrium effects.

The reduction in consumer drug expenditure generated by public pharmacies compensates for their costs. We develop a simple accounting framework to implement this comparison. First, we estimate the cost of public pharmacies using detailed data on municipal finances. We find that public pharmacies increase net expenditures of municipalities on pharmaceuticals by X%. Still, we cannot rule out that this small financial burden came at the cost of foregone increases in spending on other health goods or reduced spending on non-health goods. Second, we quantify the benefits public pharmacies provide to consumers. Combining our estimates of economic effects with summary statistics on drug expenditures and prices, we find that introducing public pharmacies in every county would reduce yearly drug expenditure by 1.5 percent or US$58 million, which is 4 percent higher than the cost of the policy.[2] Equilibrium price responses by private pharmacies are quantitatively relevant, and omitting them would lead to overestimating the reduction in expenditure by 64 percent.

[2] In addition to its economic effects, increased access to drugs could improve prescription adherence and thus health outcomes. Using data on avoidable hospitalizations and deaths, we find no evidence of such effects. This null result justifies our focus on reduced drug expenditure as a measure of benefits from public

Budget constraints and electoral incentives are crucial drivers of policy decisions (Besley and Case, 1995; Lizzeri and Persico, 2001; List and Sturm, 2006). Although we document that public pharmacies are relatively low cost and descriptive patterns suggest that mayors expected political returns, their small negative impact on a large number of people suggests that this policy might not be politically profitable. Using our field experiment, we provide suggestive evidence showing that the entry of public pharmacies increased political support for incumbent mayors. In particular, we show that awareness of the availability and attributes of a public pharmacy increased the likelihood of supporting the mayor by 6 percentage points in the local election, although point estimates are only marginally significant at conventional levels. We combine these results with our estimates of economic effects and we cannot rule out that public pharmacies have a political return that is similar to that of cash transfers (Manacorda, Miguel, and Vigorito, 2011).

Overall, we show that public pharmacies created winners and losers: Consumers who switched to public pharmacies benefited from lower prices and, those who did not, lost from higher prices. The public option did not become a financial burden because of its higher bargaining power in the input market and because private firms hold substantial market power in the wholesale and retail markets. Our paper highlights that state-owned firms could be particularly effective in other contexts in which these two conditions are also met. By doing so, we inform the long-standing question of state versus private ownership of firms and the desirability of introducing a public option into otherwise private markets. Access to a public option exists in a variety of settings, including trash collection, mail delivery, housing finance, and internet service providers in the U.S., and historically in retail gasoline stations in Canada (Petro Canada). Recent calls for the introduction of a public option in the U.S. include non-commercial banking, mortgages, and most notably healthcare.[3]

Most previous empirical work has studied public competition in the context of large programs in education (Epple and Romano, 1998; Hoxby, 2000; Dinerstein and Smith, 2021; Dinerstein, Neilson, and Otero, 2022) and health insurance (Duggan and Scott Morton, 2006; Curto, Einav, Finkelstein, Levin, and Bhattacharya, 2019). Recent work has focused on the role of state-owned firms in local markets, either directly managed by the central government, as in the case of milk stores in Mexico (Jiménez-Hernández and Seira, 2022) and branches of government-owned banks in Brazil (Fonseca and Matray, 2022), or outsourced to the private sector in the Dominican Republic and Indonesia (Busso and Galiani, 2019; Banerjee, Hanna, Kyle, Olken, and Sumarto, 2019). Relatedly, Handbury and Moshary (2021) study the price responses of grocery stores following the expansion of the national school program in the U.S. This work mostly finds that prices decrease upon increasing public competition. Our paper contributes to this literature by studying the effects of the entry of locally managed state-owned firms into local pharmaceutical markets, and by showing that public competition can potentially induce market segmentation and lead to an increase in prices by private firms.

pharmacies.

[3] See, e.g., "Why America needs a public option for mortgages" by Jeff Spross (The Week, 2017), or "There Should Be a Public Option for Everything" by Ganesh Sitaraman and Anne L. Alstott (New York Times, 2019).

This paper also contributes to a literature that studies how store entry affects local market outcomes (Basker, 2007; Hausman, 2007; Jia, 2008; Matsa, 2011; Atkin, Faber, and Gonzalez-Navarro, 2018; Arcidiacono, Ellickson, Mela, and Singleton, 2020; Bergquist and Dinerstein, 2020). The extent to which entry can generate segmentation in differentiated product oligopoly markets has been studied theoretically by Chen and Riordan (2008). Empirically, Frank and Salkever (1997) and Ward, Shimshack, Perloff, and Harris (2002) provide evidence for price increases by incumbent products upon the entry of generic drugs and private-label consumer packaged goods. We contribute to this literature by studying the consequences of entry by low-price and low-quality firms and providing evidence of market segmentation.

Our analysis of political support for incumbent mayors who opened public pharmacies is related to a large literature that studies whether and how information about politicians and policies can shape political preferences. Previous research has studied the impact of information on the candidates in an election, incumbent policies, and the prevalence of corruption (Ferraz and Finan, 2008; Gerber, Gimpel, Green, and Shaw, 2011; Chong, De La O, Karlan, and Wantchekon, 2015; Kendall, Nannicini, and Trebbi, 2015; Dias and Ferraz, 2019). Our experimental analysis differs from previous work by providing information on a specific policy directly to the people most likely to be affected by it and only a few weeks before the election.[4] More generally, we contribute to the literature by providing novel evidence of political returns to the introduction of state-owned firms in local markets.

Finally, this paper contributes to the literature that analyzes policies that aim to increase access to pharmaceuticals. Although access to affordable drugs is a first-order policy concern in low- and middle-income countries, which policies regulators should implement to achieve this goal is up for debate (UN, 2010; Pinto, Moreno-Serra, Cafagna, and Giles, 2018). Recent work examines the effects of increased competition in the retail market. Moura and Barros (2020) study the price effects of competition in the market for over-the-counter drugs, while Bennett and Yin (2019) study the price and quality effects of the entry of pharmacy chains in a market dominated by low-quality firms. Other research focuses on the effects of policies to lower drug prices, including price regulation (Dubois and Lasio, 2018; Dubois, Gandhi, and Vasserman, 2022; Mohapatra and Chatterjee, 2020; Maini and Pammolli, 2022); quality regulation (Atal, Cuesta, and Sæthre, 2022b); and public procurement (Brugués, 2020; Dubois, Lefouilli, and Straub, 2021). We provide novel evidence of how public competition in the retail market affects equilibrium market outcomes.

[4] The focus on health relates our paper to recent work on the effects of the Medicaid Expansion on voter registration and turnout (Haselswerdt, 2017; Clinton and Sances, 2018; Baicker and Finkelstein, 2019).

1.2 The Public Option in Retail Pharmaceutical Markets

Before the introduction of public pharmacies in Chile, consumers could obtain pharmaceutical drugs by buying from private pharmacies or from public health care providers. According to the 2016-2017 National Health Survey (*Encuesta Nacional de Salud*, ENS), almost 40 percent of pharmaceuticals were purchased in the private retail sector, in which there is limited insurance coverage; pharmaceuticals are the most important item of out-of-pocket health expenditures in the country (OECD, 2015; Benítez, Hernando, and Velasco, 2018).[5] The private sector is highly deregulated, as there are no market structure regulations or price controls. The three largest chains account for around 80 percent of the market share (FNE, 2019), and stores are geographically clustered in relatively rich areas (MINECON, 2013). Average profit margins in the retail sector reached 40 percent during our period of study (FNE, 2019). The wholesale market is also highly concentrated. According to data from the Economic National Prosecutor (*Fiscalía Nacional Económica*, FNE), 72 percent of off-patent medical products—defined as a unique combination of an active ingredient and a dosage—are produced by only one manufacturer, and 99 percent of those markets have an HHI above 2,500. Moreover, profit margins for manufacturers of off-patent products were 52 percent on average (FNE, 2019).[6]

The rise of public pharmacies was preceded by a collusion scandal in the pharmaceutical industry in 2008 that involved the three largest pharmacy chains in the country (Alé-Chilet, 2018). In a high-profile antitrust case, the pharmacy chains were found guilty. A left-wing mayor of a large county responded to public demands and opened the first public pharmacy in October 2015. Soon after, the popularity of the mayor boomed and dozens of other mayors from all political parties decided to open public pharmacies in the following months. By the end of 2018, 146 out of the 344 counties in the country were operating a public pharmacy. Figure 1.1 plots the number of counties with a public pharmacy over time, and Figure A.4.1 displays photos of a private and a public pharmacy.

Public pharmacies offer lower prices because they operate as nonprofit firms by law and have a cost advantage. The latter comes to a large extent from their ability to use a public intermediary that aggregates demand from public providers—most importantly, public hospitals and primary care centers—to negotiate lower prices with manufacturers. As we discuss in detail in Section 1.3 below, around two-thirds of public pharmacies purchase most of their drug supplies through the public intermediary (as opposed to directly from manufacturers). The beneficiaries of public pharmacies are determined by a combination of eligibility requirements, health conditions, and location. Most public pharmacies require that

[5] There is no broad prescription drug insurance market in Chile. Instead, there are a few disjoint programs that mostly cover drugs in the public network or for a limited set of diseases.

[6] Using a broader definition of a market that includes different dosages of the same active ingredient (ATC5), the share of single-firm markets is 54 percent. Still, 89 percent of markets have an HHI above 2,500 under that market definition.

consumers reside in the county, which is determined through a simple enrollment process that entails showing proof of residence. Also, most public pharmacies offer prescription drugs with a focus on drugs that target chronic conditions. Hence, individuals with chronic conditions are more likely to benefit. Finally, public pharmacies enter the market with a single location per county, whereas there are multiple private pharmacies in each market; this implies that for most consumers, travel costs to public pharmacies are higher than to private pharmacies.

The increasing popularity of public pharmacies has been accompanied by economic and political controversies. On the economic side, there are two main criticisms. First, that public pharmacies may be financially unsustainable and could become a burden for local governments. Second, that public pharmacies could be a form of unfair competition, particularly with respect to non-chain private pharmacies—which accounted for around 20 percent of the market, had limited buying power, and were not involved in the collusion scandal. These criticisms motivate part of our analysis, particularly the impact of public pharmacies on private sector outcomes and municipal finances.

1.3 Research Design

Data

We collected the opening dates and locations of public pharmacies. Openings span the period between October 2015 and April 2018. Figure 1.1 shows the number of openings per month and the evolution of the total number of public pharmacies operating over time. Their opening before the local election on October 23, 2016—in which most incumbent mayors were running for reelection—seemed far from a coincidence for many. The abrupt increase in openings during the months before the election is hard to explain without resorting to a political argument.

Regarding the supply of drugs by public pharmacies, we exploit detailed data on drug purchases for the 96 pharmacies that have used the public intermediary. These data include the name, molecule, dosage, amount, and price of every drug transaction by public pharmacies in 2016–2018. These data provide information on wholesale (as opposed to retail) prices, but public pharmacies charge low or no markups. While these data cover purchases through the public intermediary in detail, we have only limited data on direct purchases by public pharmacies from manufacturers. Therefore, we are unable to measure aggregate sales by public pharmacies and hence we cannot estimate the impact of their entry on aggregate sales in the market. Our limited data on direct purchases to manufacturers suggest that public pharmacies that deal with the public intermediary purchase most of their drugs through that channel.[7] Hence, we consider that the data from the public intermediary provides a fairly

[7] With the goal of measuring the relative relevance of the public intermediary as a supplier of public pharmacies, we collected additional data on public pharmacy direct purchases to manufacturers through data requests. Using data from a sample of 14 counties for which we obtained such information, we estimate that the public

accurate characterization of public pharmacies. Therefore, we use these data in Section 1.4 to describe how prices, quantities, and variety in public pharmacies compare with those in private pharmacies.

To measure outcomes for private pharmacies, we use data from IQVIA, a company that collects pharmaceutical market information worldwide. These data contain monthly local drug prices and sales for 2014-2018 collected from two sources. The four largest pharmacy chains, which account for more than 90 percent of market share, report retail prices and sales directly to IQVIA. Data for other pharmacies are collected from wholesalers.[8] IQVIA aggregates the data at the level of 66 local markets, which cover most of the country.[9] We restrict our attention to prescription drugs, which account for 93 percent of the drugs among the molecules we include in the analysis.

The Entry of Public Pharmacies

In this section, we describe entry patterns of public pharmacies and discuss how they can be exploited to study their effects. We begin with a characterization of the counties that opened a public pharmacy. We then study the timing of the entry of public pharmacies and their location within the counties in which they opened. Our results show that counties that open public pharmacies differ systematically from those that do not, but the timing of opening among those that open does not seem to be driven by observable county characteristics.

We start by comparing counties with and without public pharmacies. Columns (1)-(3) in Table 1.1 show these results. Panels A and B show that public pharmacies opened in dense high-income counties with more penetration of private health insurance, slightly better self-reported health, and a private pharmaceutical market with more pharmacies, more sales, and higher prices. In contrast, Panel C shows few differences in political variables, as measured by the previous local election of 2012.[10] If anything, counties with a public pharmacy had more candidates and were more likely to have a winner from the left wing. In sum, counties with and without public pharmacies differed significantly in terms of their pharmaceutical

intermediary accounts for around 70 percent of total purchases by public pharmacies, and is hence their main supplier. This finding motivates using the detailed data from the public intermediary in order to describe the attributes of public pharmacies.

[8] We adjust these prices for inflation using the health CPI from the National Institute of Statistics and compute prices per gram of the active ingredient to normalize them across presentations.

[9] Moreover, the data provide price and sales information at the product level for branded drugs, which identifies the laboratory, dosage, and presentation of each drug. However, for unbranded drugs it only provides price and sales information at the dosage and the presentation level, aggregated across laboratories. This is irrelevant for our analysis since we focus on price indices and aggregate sales at the molecule level.

[10] In Chile, all mayors are elected simultaneously by a simple majority rule in elections held every four years and without term limits until 2020. To measure local political outcomes, we use county-level information about candidates, parties, coalitions, and votes for each candidate in the 2012 and 2016 local elections from the Electoral Service. The 2012 election allows us to characterize the political equilibrium before the opening of public pharmacies.

market and socioeconomic characteristics but were relatively more similar in their political characteristics.

To examine entry timing systematically, we ranked all public pharmacies by their entry date and estimated an ordered logit model of this ranking on all variables in Table 1.1. Column (4) in the table presents the results. Pharmacies that opened earlier entered counties with more population and were more likely to have left-wing mayors, but entry timing is otherwise uncorrelated with the characteristics of the pharmaceutical market, socioeconomic attributes, or electoral competition in the previous election. Instead, anecdotal evidence suggests that unexpected delays in sanitary permits explain why some pharmacies opened after the election. We rely on these results to exploit the timing of entry as exogenous variation.

Finally, we document that mayors opened public pharmacies near existing private pharmacies, which provides a unique opportunity to study the impact of the public option in an existing market. To describe their location choices, we geocoded all private pharmacies in the country and assigned them to geographic cells of 600×600 meters. We then estimated cross-sectional cell-level regressions using data from counties with a public pharmacy. The dependent variable is an indicator for a cell that has a public pharmacy, and explanatory variables include the number of private pharmacies, the number of schools as a proxy for population, and county-level fixed effects. Table A.4.1 displays the results. Estimates reveal that public pharmacies opened in populated areas where private pharmacies were already operating. The maps in Figure 1.2 provide visual examples of the entry decision in six counties spread across the country.

1.4 The Economic Effects of Public Pharmacies

Evidence on Prices and Quality of Public Pharmacies

When public pharmacies opened, consumers gained access to a new alternative in their choice set, which differed from available options along several dimensions. We describe the basic attributes of public pharmacies by using transaction-level data on all purchases by public pharmacies from the public intermediary in 2016–2018. The public intermediary was the main supplier of drugs for the 96 counties that sourced through it, as discussed in Section 1.3.

The most salient and advertised difference was related to drug prices. Using a set of exactly matched drugs that are sold in both public and private pharmacies, we study price differences across public and private pharmacies. In Figure 1.3, Panel (a) shows that almost all drugs are sold at lower prices in the former and that the relative price difference is, on average, between 64 and 68 percent depending on the margin public pharmacies charge over purchase costs from the public intermediary. These large price differences suggest that consumers should, in principle, switch to public pharmacies in the local markets in which they open.

Two leading reasons for these price differences are public pharmacies' higher bargaining power in the input market—coupled with a concentrated input market—and the substantial market power of private retailers downstream, both of which we discussed in Section 1.2. In A.1, we formalize these arguments by developing a model of the vertical chain that captures the main features of our setting—namely, that (i) producers and retailers are able to exercise market power, (ii) state-owned firms differ from private firms by having greater bargaining power upstream, and (iii) state-owned firms do not maximize profits but rather total surplus. We show that under mild assumptions regarding the demand curve, downstream prices are lower when retailers have more bargaining power upstream and when retailers place a higher weight on consumer welfare relative to profits.

Consumers trade off lower prices with the lower quality of public pharmacies. The fact that public pharmacies enter with a single store in each county implies that most consumers have multiple private pharmacies closer to their homes. Using data on voter home addresses from the Electoral Registry, and the locations of public and private pharmacies, we calculate distances between households and every pharmacy in the county. The average (median) individual has 20 (12) private pharmacies located closer than the public pharmacy in their county. Panel (b) in Figure 1.3 shows that the distributions of distance to the closest private pharmacy and public pharmacy differ markedly: The average distance to the closest private pharmacy is 1.1 kilometers—less than half of that to the public pharmacy. These facts imply that shopping at public pharmacies entails higher travel costs than shopping at private pharmacies. Moreover, public pharmacies offer less product variety. Panel (c) in Figure 1.3 shows that the average number of products per molecule-county is 2.2, and that 70 percent of molecule-counties offer 3 varieties or fewer, while the average number of varieties in private pharmacies is 15.2.[11] To the extent that consumers value product variety, these patterns imply that public pharmacies are less convenient than private pharmacies. The longer waiting times and limited opening hours already described in Section 1.2 further exacerbate the relatively low quality of public pharmacies.

The relevance of public pharmacies has grown over time, which demonstrates that at least some consumers value lower drug prices relative to lower convenience enough to switch to public pharmacies. Panel (d) in Figure 1.3 shows that their average market share across molecules and counties reached around 4 percent by the end of 2018. Of course, it is unclear whether sales by public pharmacies have decreased sales by private pharmacies or simply expanded market size. To inform this margin, we estimate the effects of public pharmacies on private pharmacy sales.

[11] Relatedly, public pharmacies are more likely to offer only generic drugs or only branded drugs within a molecule: This is the case for 72 percent of molecule-counties at public pharmacies, but for only 36 percent at private pharmacies.

Equilibrium Effects on Prices and Sales by Private Pharmacies

Public pharmacies may induce consumers to substitute away from private pharmacies.[12] Moreover, the competitive pressure from public pharmacies may induce private pharmacies to adjust prices. In this section, we estimate the effects of the entry of public pharmacies on prices and sales by private pharmacies.

Theoretically, the effects of entry on incumbent firm prices are ambiguous. Chen and Riordan (2008) study the conditions under which entry leads to increases or decreases in prices. Their analysis shows that these effects depend on the magnitudes of two effects of entry on the incumbent's pricing incentives. First, entry has a *market share effect*, which depends on the extent to which the incumbent loses demand upon entry due to substitution. The more demand the entrant takes away from the incumbent, the stronger the incentives for the incumbent to decrease prices in response to entry. Second, entry has a *price sensitivity effect*, which depends on how the slope of the incumbent's residual demand curve changes after entry. The steeper the demand curve after entry relative to before entry, the lower the extent of substitution away from the incumbent upon entry, and therefore the stronger its incentive to increase prices upon entry. Overall, the incumbent's price will increase whenever the price sensitivity effect dominates the market share effect and vice versa. Which effect dominates depends on the distribution of consumer preferences and on the attributes of the firms. To further develop intuition for the conditions under which private pharmacy prices may decrease or increase upon the entry of public pharmacies, we develop a model based on Chen and Riordan (2008) in A.3. We then implement illustrative simulations that we employ to discuss our results.

Event study evidence. We start by exploiting the staggered entry of public pharmacies in an event study framework. For this analysis, we use IQVIA data on drug prices and sales across local markets. A challenge in combining data on the entry of public pharmacies with data from IQVIA is that the level of geographic aggregation of the latter markets is in some cases larger than counties, which is the level at which public pharmacies operate. To tackle this issue, we estimate a stacked event study regression.[13] Whenever a market has more than one event, we create as many copies of the data as the number of events. We stack the

[12] As part of this research, we designed and implemented an informational field experiment to study the impacts of public pharmacies. In the experiment, we randomly provided information about public pharmacies to individuals buying pharmaceuticals in private pharmacies. In this paper, we use the experiment to estimate the impact of public pharmacies on support for incumbent mayors who opened these. We provide more details in Section 1.6. However, we also collected data on consumer shopping behavior both before and two months after the intervention, to study whether consumers in the pharmaceutical market switched from private to public pharmacies. Overall, consumers learned about the low-price and low-quality of public pharmacies after the intervention, and to some extent reported either having used or planning to use the public pharmacy. We discuss these findings in A.2.

[13] This approach has been adopted in recent work that estimates event study models in settings with multiple events per unit (see, e.g., Lafortune, Rothstein, and Schanzenbach 2018; Cengiz, Dube, Lindner, and Zipperer 2019b).

copies in a dataset and use the entry of public pharmacies to all counties within a market as events. Figure A.4.2 shows the distribution of the number of events per market.

The main specification we estimate is given by:

$$y_{mlgt} = \sum_{k=-12}^{18} \beta_k D_{lgt}^k + \lambda_{mt} + \theta_{mlg} + \varepsilon_{mlgt}, \tag{1.1}$$

where g indexes entry events within a market. The dependent variable y_{mlgt} is either the log of drug prices or the log of drug sales for molecule m in local market l in month t.[14] Our interest is in the coefficients β_k on the dummies $D_{lgt}^k = 1\{t = e_{lg} + k\}$, which indicate whether a month t is exactly k months after event time e_{lg} for event g in local market l. We normalize $\beta_{k=-1} = 0$, so we interpret all coefficients β_k as the effect of a public pharmacy's opening on the dependent variable exactly k months after its entry. The specification also includes molecule-month fixed effects λ_{mt} to account for time-varying unobservables at the level of molecules, and molecule-market-event fixed effects θ_{mlg} to account for persistent differences in market conditions across markets. Standard errors are clustered at molecule-market level.[15]

The entry of public pharmacies had meaningful effects on private pharmacies. Panels (a) and (b) in Figure 1.4 present the results for sales and prices, respectively. Drug sales by private pharmacies decrease after a public pharmacy enters a market. Our estimates imply that 18 months after the entry of a public pharmacy, private pharmacies in that market sell around 4 percent less. Furthermore, 18 months after the entry of a public pharmacy, drug prices in private pharmacies increase by 1 percent. Both effects increase over time, which suggests that public pharmacies evolve in terms of enrolling more consumers and possibly improving their product offerings and convenience.[16]

The main threat to identification of the effect of public pharmacies is reverse causality; unobserved determinants of sales and prices in the private sector may drive the entry of public pharmacies. In that case, β_k would confound the causal effect of public pharmacies on

[14] We define the market-level price as the share-weighted average of log prices:

$$\hat{P}_{mlt} = \sum_{i \in \mathcal{I}_{ml}} w_{il0} P_{ilt},$$

where $\mathcal{I}_{ml}$ is the set of drugs of molecule m in local market l, P_{ilt} is the log price per gram of product i in period t and market l, and w_{il0} denotes the share of sales of drug i in market l in 2014. Because these weights are constant, changes in the index are driven by changes in prices and not by changes in market shares or market structure. This price index has been used in previous work studying retail drug pricing (e.g., Atal et al., 2022b). For sales, we use the residuals from the projection of the outcome variable on month-of-the-year fixed effects by molecule-market to account for seasonality that is specific to sales in some markets (e.g., due to tourism in the summer).

[15] We use a balanced sample of markets in event time and include never-treated markets to pin down the linear component of pre-trends (Borusyak, Jaravel, and Spiess, 2022b). Moreover, we fully saturate the model and report results for event dummies 12 months before and 18 months after the event.

[16] An additional margin of response for private pharmacies would be to adjust product variety. We estimate equation (1.2) using the number of varieties offered as the dependent variable, and find no evidence of responses along that margin.

private market outcomes with trends in outcomes that cause the entry of public pharmacies.[17] Reassuringly, the lack of pre-trends in both sales and prices leading up to the entry of public pharmacies suggests that reverse causality and strategic considerations do not play a significant role in our setting.[18]

Another concern relates to multiple public pharmacy entries within a market, which could potentially turn the treatment effect of a previous public pharmacy entry into a pre-trend for the subsequent entry. This concern is muted in our context because the majority of markets experience 1 or 2 events and most subsequent entry occurs within 1 or 2 months of each other, as shown by Figure A.4.2. To assess the importance of this issue in our setting, we do two robustness checks. First, we redefine the event as the $first$ entry of a public pharmacy, in which case this type of pre-trend is absent by definition. The results under that treatment definition are essentially the same as those in our main specification, as shown by Figure A.4.3. Second, we restrict the estimating sample to markets with a single event or multiple events separated by less than 1 month. The results for this sample track closely those from our main sample, as shown by Figure A.4.4.

Exposure difference-in-differences design. We complement the event study design with a regression analysis that relates market-level outcomes to the share of the population in each market that has access to a public pharmacy at each point in time. The advantage of this design is that it exploits all the variation in the timing of entry of public pharmacies as well as the heterogeneous exposure of markets to public pharmacies. We then employ this design to develop a heterogeneity analysis for the effects of public pharmacies.

We define treatment intensity E_{lt} as the share of the population in market l with access to a public pharmacy at time t, and estimate the following specification:

$$y_{mlt} = \lambda_{mt} + \theta_{ml} + \beta^{\text{jump}} E_{lt} + \beta^{\text{phase in}} E_{lt}(t - t_e^* + 1) + \varepsilon_{mlt}, \tag{1.2}$$

where $E_{lt} = 0 \ \forall \ t < t_e^*$. This functional form is motivated by the patterns of the treatment effects we estimate in our event study analysis in Figure 1.4. The parameter β^{jump} is a mean shift in outcome y_{mlt} after the adoption of a public pharmacy. Since results from the event study specification imply that the impact on sales and prices evolves over time, we allow for a trend break, $\beta^{\text{phase in}}$. We include event-time dummies as controls for all periods before $k = -12$ and after $k = 18$ in treated markets, for comparability with the event study results. Our main parameter of interest is the effect of the public pharmacy 18 months after its entry, which we calculate as $\bar{E}_{18} \times [\beta^{\text{jump}} + (18+1)\beta^{\text{phase in}}]$. The term $\bar{E}_{18}$ is the average exposure

[17] Strategic entry is an identification threat for reduced-form models for the effects of firm entry as equation (1.1), but it is not a relevant concern in our context. Public pharmacies' business model differs from private pharmacies' since they operate as nonprofit firms.

[18] As an additional piece of supporting evidence, in column (4) of Table 1.1 we study the order of entry of public pharmacies using an ordered logit regression of entry on market and political covariates. The results show that the timing of entry is uncorrelated with covariates associated with the supply and demand of drugs.

to a public pharmacy across markets 18 months after the entry of the first pharmacy in the market.

For ease of exposition, we present the results of the main parameter of interest in Table 1.2 and report the underlying estimates β^{jump} and $\beta^{\text{phase in}}$ in Table A.4.2. Columns (1) and (2) in Table 1.2 present estimates for sales and prices, respectively. Panel A shows that the entry of public pharmacies decreases drug sales by private pharmacies by 3.8 percent and increases drug prices by private pharmacies by 1 percent 18 months after their introduction. Reassuringly, these magnitudes are close to the estimates we obtain at the end of the time window in the event studies in Figure 1.4. To put the magnitude of this estimate in context, the average coefficient of variation of drug prices across drugs and local markets is 0.08. Hence, our estimates imply that drug prices at private pharmacy prices increase by around 12.5 percent of a (relative) standard deviation after the entry of a public pharmacy.[19]

Heterogeneity analysis. The remaining panels in Table 1.2 present a heterogeneity analysis. The characteristics of the context motivated us to focus on three margins. First, public pharmacies specialize in selling drugs for chronic conditions and thus we expect a larger impact on these drugs. Column (1) in Panel B shows that sales of chronic drugs decrease by 4.5 percent, which is 61 percent more than the 2.8 percent decrease in non-chronic drugs (p-value<0.01).[20] In contrast, column (2) in Panel B shows similar price increases for both types of molecules. Second, we have emphasized quality differences across public and private pharmacies. We proxy relative quality by the ratio of drug variety within each molecule in public pharmacies relative to private pharmacies in each market.[21] Column (1) in Panel C shows that the impact is larger in markets in which the public pharmacy has a richer variety of products within each molecule (p-value 0.02). Column (2) in Panel C reveals larger price responses in markets in which public pharmacies offer less variety of products within a molecule (p-value<0.01). Finally, we consider whether the spatial distribution of private pharmacies matters for the impacts of public pharmacies. We expect that the closer public pharmacies locate to private pharmacies, the larger the decrease in private pharmacy sales. Column (1) in Panel D presents heterogeneous effects along this dimension and confirms this intuition (p-value 0.05).[22]

[19] The extent of price variation in our data is somewhat higher than roughly comparable measures for within-chain pricing reported by Adams and Williams (2019) and DellaVigna and Gentzkow (2019) for construction materials and consumer-packaged goods in the U.S, respectively. This price variation is consistent with our ability to estimate price effects in this setting. Results available from the authors.

[20] We observe 102 chronic molecules and 74 non-chronic molecules. This finding is consistent with our experimental evidence showing that households with members with chronic conditions react more strongly to the availability of public pharmacies in terms of shopping behavior. We discuss experimental results in A.2.

[21] We define high (low) variety as observations above (below) the median of the ratio between the number of distinct products within molecule and market offered by the public pharmacy and those by private pharmacies.

[22] To split the sample in two, we use the average number of public pharmacies operating within 400 meters of private pharmacies. For consistency, we only consider private pharmacies that appear in our data for private pharmacy outcomes. These results need to be interpreted with caution as public pharmacies mostly locate

Discussion

The entry of public pharmacies had equilibrium effects on private pharmacies. As expected, due to the lower prices offered by public pharmacies, some consumers substituted away from private pharmacies and drug sales in the latter decreased. While increased competition could have induced private pharmacies to reduce drug prices, we find that private pharmacies instead increased prices. This response is consistent with the price sensitivity effect of entry dominating the market share effect of entry. In particular, while some consumers switched to public pharmacies upon their entry, it must be that they had a relatively low willingness to pay for private pharmacies, which led to the residual demand for private pharmacies to become steeper. The increase in private pharmacy prices we estimate implies that the upward pricing pressure from the latter was larger than the downward pricing pressure from overall substitution toward public pharmacies.[23][24]

The sales response to the entry of public pharmacies may seem small, given the magnitude of the price differences between public and private pharmacies. Our interpretation is that product differentiation plays a role in mediating this response. As documented above, public pharmacies are less convenient than private pharmacies in terms of waiting times, opening hours, product variety, and travel distance. The lack of a stronger response suggests that a sizable share of consumers value those attributes enough to not substitute toward public pharmacies on the basis of lower prices. Higher-quality public pharmacies would have likely led to stronger equilibrium responses.[25] Second, our event study results in Figure 1.4 show that both quantity and price effects increase over time, which suggests that the full effects may be larger once the market settles into a new equilibrium.

The substitution away from private pharmacies we estimate is consistent with findings

nearby private pharmacies and information about how distance affects pharmacy choice is lacking.

[23] In our model in A.3, we show that a key condition under which private pharmacy prices are more likely to increase is a negative correlation in consumer willingness to pay for public and private pharmacies, such that consumers who have a high valuation for private pharmacies also have a low valuation for public pharmacies. This negative correlation implies that consumers who substitute away from the private pharmacy upon entry are those with low willingness to pay for the private pharmacy—and thus the most price sensitive—which leads to the residual demand curve of the public pharmacy's being steeper after entry. In addition, there must be enough heterogeneity in willingness to pay across consumers, as otherwise there is no scope for increasing prices substantially. Figure A.4.5 shows simulation results that demonstrate that the direction of the price effects of entry indeed depends on these parameters of the distribution of consumer preferences.

[24] Caves, Whinston, Hurwitz, Pakes, and Temin (1991) and Frank and Salkever (1997) document a similar pattern of market segmentation in pharmaceuticals, in which innovator drugs that become off-patent do not decrease but rather *increase* their prices after generic entry. This fact is known in the literature on competition in pharmaceutical markets as the "generic paradox."

[25] We illustrate the role of vertical differentiation between private and public pharmacies using our model in A.3. Our model simulations show that vertical differentiation indeed influences the extent to which the entry of public pharmacies affects private pharmacy prices, and market share depends on vertical differentiation. Panel A in Figure A.4.6 shows that the extent of business stealing by an entrant decreases substantially as the quality of the entrant relative to the incumbent decreases. Moreover, Panel B in Figure A.4.6 shows that the incumbent in the market is able to sustain higher prices when the quality of the entrant relative to the incumbent is lower.

in related work by Busso and Galiani (2019) and Jiménez-Hernández and Seira (2022) in different contexts. However, they find a price decrease among private firms as opposed to a price increase. Our results highlight the fact that the price effects of public competition will depend on underlying consumer preferences and firm attributes.

1.5 The Benefits and Costs of Public Pharmacies

This section discusses the relative efficiency of state-owned firms. First, we estimate the cost of public pharmacies by exploiting data on municipal finance to study the effects of introducing public pharmacies on spending and revenue on health and non-health services. Second, we assess whether public pharmacies have any health effects on consumers as measured by avoidable hospitalizations. Finally, we develop a simple framework that exploits our estimates of the price and quantity effects of public pharmacies to estimate how consumer drug expenditure decreases as a result of public pharmacies, and compare it with our cost estimates.

Municipal Finance and the Cost of Public Pharmacies

Given that public pharmacies were created by local governments that manage multiple other local services, it is important to identify whether they are economically sustainable or represent a financial burden that may crowd out other services. To study this margin, we exploit administrative data from municipal finances to estimate the financial impacts of public pharmacies.[26]

For this analysis, we estimate the following regression:

$$y_{ct} = \theta_c + \lambda_t + \pi^{\text{jump}} PP_{ct} + \pi^{\text{phase in}} PP_{ct}(t - t_e^* + 1) + \varepsilon_{ct}, \quad (1.3)$$

where y_{ct} is a financial outcome in county c and year t (e.g., spending on health services), PP_{ct} indicates the share of the year with a public pharmacy in county c, and $t - t_e^*$ measures the number of years since the opening of the public pharmacy. The specification includes county fixed effects θ_c and year fixed effects λ_t. Similar to our specification for private market outcomes in equation (1.2), the parameter π^{jump} captures a mean shift in the dependent variable after treatment, whereas $\pi^{\text{phase in}}$ captures a trend break. In terms of data, we observe annual county spending and revenue for 2013–2019. Both spending and revenue have subcategories we aggregate into health and non-health categories. To ease comparison

[26]The data come from the National System of Municipal Information (*Sistema Nacional de Información Municipal*, SINIM). Counties spend resources on transportation, public education, public health, culture, and sports, among others (Law 18695). Approximately 90 percent of their budget comes from county revenues (property and vehicle tax receipts) and other resources correspond to monetary transfers from the central government.

across counties, we use the log spending and revenue per capita as dependent variables in this analysis.[27]

Table 1.3 presents our main results and Table A.4.3 presents coefficient estimates of equation (1.3). The main result is the effect of public pharmacies after 18 months of operation (1.5 years), which we compute as $\pi^{\text{jump}} + (1.5 + 1) \times \pi^{\text{phase in}}$. The results deliver three main messages. First, 18 months after the entry of public pharmacies, we observe an increase of 4.1 percent in health spending in column (1), which is partially compensated for by an increase in health revenue of 2.7 percent in column (2). The difference between these effects is statistically significant (p-value 0.036). Second, the impact of public pharmacies on non-health services in columns (3) and (4) is imprecisely estimated and we cannot rule out a decrease of a magnitude similar to the increase in health services. Third, in terms of overall municipal finance, our point estimates in columns (5) and (6) imply that spending increases more than revenue, although those coefficients are again not statistically significant. Taken together, the point estimates in the last two columns suggest that public pharmacies induced, if any, only a small and statistically insignificant increase in the overall municipal deficit.[28]

These estimates allow us to compute the average cost of introducing a public pharmacy. A public pharmacy's profits depend on the markup they charge on drugs if any, and any initial investment and operating cost it incurs. The fact that public pharmacies induce a deficit implies that they set prices below average cost. The average spending and revenue per capita are \$695.68 and \$730.15 and the average county in the country has a population of 51,781. Combining these basic statistics with our point estimates in columns (5) and (6) of Table 1.3, we calculate that after 18 months of operation the annual loss for a public pharmacy in the average county is \$162,266.[29] The next sections compare this cost estimate with the estimated benefits of public pharmacies for consumers.

Lack of Health Effects of Public Pharmacies

Increased access to pharmaceutical drugs could benefit individuals through health improvements. For instance, such effects could operate through improved adherence to prescription drugs for individuals with chronic diseases due to lower prices and increased access (Cutler and Everett, 2010). However, in our setting we do not observe individual-level prescriptions and drug purchases. Instead, we focus on avoidable hospitalizations associated with chronic diseases, which would likely have not occurred under appropriate disease management. This variable has been employed previously in the literature (e.g., Layton, Maestas,

[27] Some counties, which account for 7 percent of the sample, do not report the breakdown of their accounts for health and non-health services. To obtain a uniform sample across dependent variables, we drop those observations.

[28] Figure A.4.7 displays corresponding event study estimates and provides reassuring evidence regarding the trends in these outcomes leading up to the entry of public pharmacies.

[29] Articles from local newspapers that disclose public pharmacy non-drug costs place the yearly cost of running them at between \$85,000 and \$125,000, which likely provide a lower bound for total operating costs and are in line with our estimates (see, e.g., Araucanía Cuenta 2016; El Austral 2017; Clave9 2017; Diario Concepción 2017; Diario Financiero 2022).

Prinz, and Vabson, 2019). The fact that public pharmacies were oriented toward individuals with chronic diseases makes this variable particularly suitable. We would interpret a decrease in avoidable hospitalizations after the entry of a public pharmacy as a signal that the pharmacy increased drug access and, in consequence, adherence by individuals with chronic diseases.

For this analysis, we estimate equation (1.3) using avoidable hospitalizations as the dependent variable. We exploit data on monthly hospitalizations for 2013–2019 from the Ministry of Health (DEIS, 2019), which cover the number of hospitalizations, days of hospitalization, number of surgeries, and number of deaths per diagnosis across all hospitals in the country. The number of hospitalizations captures only the volume of these events, whereas hospitalization days, surgeries, and deaths capture their severity. To focus on the subset of diagnoses for which hospitalizations are considered avoidable, we follow the Prevention Quality Indicators in AHRQ (2019), which lists all diagnosis codes (ICD-10) for avoidable admissions associated with asthma, chronic obstructive pulmonary disease, diabetes, and hypertension. We restrict our sample of hospitalizations for this analysis to these diagnoses. We normalize these variables by population and measure them per 100,000 inhabitants.

Our estimates suggest that public pharmacies did not improve health outcomes, at least in the short period of time we are able to examine. Table 1.4 presents our main results and Table A.4.4 presents coefficient estimates of equation (1.3). For each outcome, we show results for all individuals and for those under public insurance (*Fondo Nacional de Salud*, FONASA), who on average have lower income and are more likely to benefit from a public pharmacy. Across all outcomes and samples, we find no statistically significant effect of the entry of a public pharmacy to a local market after 18 months. That said, our estimates are not precise enough to rule out effects that could be quantitatively meaningful. In particular, our estimates can reject at the 5 percent level reductions of 2.43 hospitalizations, 21.15 hospitalization days, 0.23 surgeries, and 0.07 deaths per 100,000 inhabitants as the effect of public pharmacies, which are equivalent to reductions of between 10 percent and 13 percent in these outcomes relative to their baseline levels.[30][31]

Overall, our interpretation of these results is that public pharmacies did not affect access to drugs to an extent such that adherence improved enough as to reduce avoidable hospitalizations. It is important to note that the lack of a health effect is likely to be mediated by contextual factors such as the elasticity of demand and access to health services, among others. Regardless, these results suggest that if public pharmacies had any market-creation effect, it was small, and most of the effect was through business stealing from private pharmacies.

[30] Figure A.4.8 shows the results of an event study version of equation (1.3). For all outcomes and samples, we again find no evidence that public pharmacies affected health outcomes. Reassuringly, these results show a lack of differential trends across counties leading up to the entry of public pharmacies, which provides evidence against reverse causality.

[31] An additional analysis of school attendance and sick leaves—arguably related to the health of children and the working population—also suggests a null impact of public pharmacies in the short run. See Table A.4.5 and Figure A.4.9.

Comparing Costs and Benefits

In this section, we use our previous results to compare the benefits and costs of public pharmacies. Our measure of benefits from public pharmacies focuses on reduced expenditure in drugs for consumers, given that we find no evidence of health effects. We develop a simple accounting framework to estimate effects on consumer expenditure by combining our results on economic effects from Section 1.4 with basic statistics from the market.

Let r denote private pharmacies and u denote the public pharmacy. Moreover, let $t = 0$ indicate the period before entry of the public pharmacy and $t = 1$ the period after its entry. Using this notation, total consumer expenditure in period t is given by $e_t = M_t(s_t^r p_t^r + s_t^u p_t^u)$, where M_t is the amount of drugs consumers need; s_t^r and s_t^u are market shares of the private and the public pharmacy, respectively; and p_t^r and p_t^u are composite drug prices at each of them. We impose two assumptions. First, we assume that the market size remains constant over time, such that $M_t = M$ for $t = 0, 1$. Second, given that we are unable to estimate aggregate effects on drug quantity with the available data, we rule out such effects and impose $s_t^r + s_t^u = 1$ for $t = 0, 1$.

The object of interest is the change in drug expenditure upon entry of the public pharmacy:

$$\Delta e \quad = \quad M(s_1^r p_1^r + s_1^u p_1^u) - M(s_0^r p_0^r + s_0^u p_0^u),$$

which we can rearrange to be a function of our estimates and data. First, note that $s_0^r = 1$ and $s_0^u = 0$ by definition. Second, we use our estimates of effects on private pharmacies from Section 1.4 to express the sales and prices of private pharmacies after the entry of the public pharmacy as $s_r^1 = (1 - \beta_s)s_0^r$ and $p_r^1 = (1 + \beta_p)p_0^r$, respectively. Finally, we use results from Section 1.4 on price differences between public and private pharmacies to express public pharmacy prices as $p_u^1 = \phi_1^u p_r^1$, where ϕ_1^u is the average discount public pharmacies offer relative to private pharmacies. After replacing and rearranging, we get:

$$\Delta e \quad = \quad \underbrace{Mp_0^r}_{\substack{\text{Baseline} \\ \text{expenditure}}} \times \Big[\underbrace{(1 - \beta_s)(1 + \beta_p) - 1}_{\substack{\Delta \text{ expenditure} \\ \text{in private pharmacies}}} + \underbrace{\beta_s \phi_1^u (1 + \beta_p)}_{\substack{\Delta \text{ expenditure} \\ \text{in public pharmacy}}}\Big].$$

To measure the change in drug expenditure, we proceed as follows. We measure baseline expenditure using data from the 2016 National Household Spending Survey (*Encuesta de Presupuestos Familiares* EPF) which states that the average yearly drug expenditures were \$213.4. Furthermore, our estimates from Section 1.4 imply that $\beta_s = 0.038$ and $\beta_p = 0.010$. Finally, we know from Section 1.4 that public pharmacies set prices at an average of $\phi_1^u = 0.34$ of private pharmacy prices.

The average consumer saves \$3.3 per year, according to these estimates. This average masks substantial heterogeneity: Those who stayed at private pharmacies increased their annual spending by \$2.1, whereas those who switched to the public pharmacy reduced theirs by \$140.1. A population of particular interest is consumers with chronic conditions, who are the main target of public pharmacies and account for 22 percent of the population, according to the 2016–2017 ENS. Our estimates imply that these consumers decreased their yearly expenditure by an average of $16.4. Of them, those who stayed with private pharmacies increased their yearly expenditure by $8.2, whereas those who switched decreased it by $537.3. To put these numbers in context, the median monthly wage among working-age individuals in 2017 was around $670. Adding across consumers, these estimates imply that consumers in the average county decreased their aggregate spending by $171,166 per year. If all counties in the country introduced public pharmacies, aggregate spending would decrease by $57.66 million per year—equivalent to 1.53 percent of total expenditure according to the EPF. Accounting for equilibrium price responses by private pharmacies is quantitatively relevant; omitting them would lead to overestimating the reduction in expenditure by 64 percent.

Our estimates imply that consumer benefits in terms of reduced drug expenditure on inframarginal units are 4.4 percent higher than the cost of public pharmacies a year and a half after their entry. Public pharmacies achieve reductions in consumer expenditure higher than their costs for two reasons: public pharmacies hold a cost advantage relative to private pharmacies when purchasing from manufacturers, and private pharmacies hold substantial market power in the retail market (FNE, 2019). Public pharmacies thus address two salient market failures in this industry. Because of this, the introduction of a state-owned firm likely performs better than an alternative policy of subsidizing drug purchases. In this simple framework, the cost of a subsidy is the reduction in drug expenditure, and is thus higher than that of the public pharmacy, according to our estimates. This is because subsidies are able to reduce drug expenditure, but do not address market power in the private market and therefore must incur a higher cost to achieve the same effects as the public pharmacy.[32]

Of course, this is not a full welfare analysis. On the one hand, we do not account for potential market expansion effects, which implies that we may underestimate the benefits of public pharmacies. On the other hand, we do not account for consumer valuation of the relative convenience of private and public pharmacies. The fact that relatively few consumers switch despite the large potential savings for switchers suggests that the valuation of these non-price pharmacy attributes is high.[33] A richer model of consumer demand and pharmacy pricing is needed to conduct such an analysis.[34]

[32] Enriching the framework to account for aggregate effects would exacerbate the extent to which state-owned firms outperform subsidies since subsidies would in that case induce an additional deadweight loss.

[33] To provide a lower bound on the relative inconvenience of public pharmacies, we estimated the cost of additional travel time to public pharmacies. To do so, we combined standard assumptions from the transportation literature with data on (i) the spatial distribution of households, private pharmacies, and public pharmacies, and (ii) the distribution of hourly wages. We find that an individual with an average hourly wage has an average annual cost of additional travel time to public pharmacies of $13.9, with 25*th* and 75*th* percentiles of $2.3 and $21, which are well below our estimates of average savings for switchers. These patterns suggest that while their inconvenient locations may indeed contribute to the low switching rate to public pharmacies, other differences between public and private pharmacies play a relevant role as well. Calculations are available from the authors.

[34] Other unmeasured welfare effects include potential decreases in incentives for R&D. However, we believe

1.6 Political Returns of Public Pharmacies

Budget constraints and electoral incentives are crucial drivers of policy decisions (Besley and Case, 1995; Lizzeri and Persico, 2001; List and Sturm, 2006). The small negative impact on a large number of consumers suggests that the public option might not be politically profitable. This section uses an informational field experiment, along with self-reported voting behavior, to estimate the causal effect of the awareness of public pharmacies among consumers in the pharmaceutical market on political support for the incumbent who opened the pharmacy.

The Field Experiment

We designed a field experiment to study whether the availability of public pharmacies affected consumers. To induce variation in awareness of the public pharmacy within local markets, we implemented an informational intervention. The decision to provide information was based on a survey we conducted before the experiment, which revealed that consumers were only partially informed along two dimensions. First, some households were unaware of the existence of a public pharmacy in their county. Second, even when households knew about the pharmacy, they were not perfectly informed about the lower prices and other attributes. The existence of imperfect information provides us with a unique opportunity to randomly expose consumers to public pharmacies using our experiment, and thus to measure individual responses to them.

The treatment consisted of an informational flyer, displayed in Figure A.4.10. It provided information about the presence of a public pharmacy in the county and stated that it offered lower prices but longer waiting times than private pharmacies. Also, it included the pharmacy's location, contact information, opening hours, and eligibility requirements. We delivered the flyer to consumers exiting private pharmacies in the 20 counties with public pharmacies in Santiago, displayed in Figure A.4.11. The information was tailored to each county.

In terms of recruitment, enumerators approached consumers leaving a private pharmacy in each county and assessed their eligibility. Eligible participants were those who (i) lived and were registered to vote in the county, (ii) had purchased a prescription drug, and (iii) were not registered with the public pharmacy. To incentivize participation, everyone who responded to the 5-minute survey automatically entered a lottery for a television set. Overall, 1,855 individuals were approached and 826 enrolled in the study. The baseline survey collected information on awareness of public pharmacies and their attributes, intention to vote for the incumbent mayor in the upcoming election, age, education, and access to the internet, among others. When the survey was completed, participants were randomly assigned to treatment and control groups. The enumerator only learned the assignment of the individual after completing the survey. We conducted this survey between October 12 and 20, 2016,

that this effect is likely small given the Chilean market represents only a small share of the revenues of the pharmaceutical companies doing R&D.

right before the local elections. Figure A.4.12 summarizes the timeline of the events in the experiment.

Two months after the baseline survey, we conducted a follow-up survey to measure the same variables as in the baseline. We also collected information about their relationship with the public pharmacy in their county. We conducted this survey by phone and were able to complete the survey for 514 participants—almost two-thirds of the sample.[35][36]

Table A.4.8 compares both groups at baseline. Participants are on average 45 years old and 61 percent of them are female. More than 60 percent work, and most use the internet frequently. Half of the participants planned to vote for the incumbent and almost three out of four reported having participated in the previous election. Slightly less than 70 percent knew about the existence of a public pharmacy. As expected, column (4) shows that almost all variables are balanced across groups. The exception is awareness of the public pharmacy, which we control for in the analysis.

Experimental Results

Table 1.5 presents results from estimating equation (A.3) for political outcomes. Columns (1) and (4) study self-reported voting behavior. As many as 28 and 26 percent of the control group individuals reported voting for the incumbent mayor and incumbent party, respectively. The reported vote increases by approximately 6 percentage points for the treatment group in both cases. While these point estimates are large in magnitude, they are not statistically significant at conventional levels, with p-values of 0.21 and 0.12. To increase the precision of the analysis, columns (2) and (5) control for the intention to vote for the mayor at baseline along other covariates, and include county fixed effects. Treatment effects using this specification remain similar in magnitude but are indeed more precise, with p-values of 0.06 and 0.11.[37]

Effects on voting behavior are concentrated among individuals from households with members with chronic conditions. Columns (3) and (6) examine these patterns of heterogeneity. Households with someone with a chronic condition report having voted 8 percentage points more for the incumbent, larger than the 2-7 percentage points higher vote share among treated households without a chronic condition. Although the small sample prevents us from

[35] Table A.4.6-A shows that attrition was higher among younger participants, males, with higher support for the incumbent, less turnout in the last election, and less knowledge of the public pharmacy. While this changes the sample composition and decreases the statistical power of the experiment, it does not necessarily threaten its internal validity. Table A.4.6-B shows that all variables remain balanced across groups among non-attriters.

[36] The survey also verified the delivery of the treatment. Table A.4.7 shows that treated individuals acknowledged receiving information more often than those in the control group, and recalled public pharmacies' being the core of the information content almost twice as often as the latter.

[37] To account for the effects of attrition, Table 1.5 presents Lee bounds. The lower bound is positive but not statistically significant and the upper bound is positive and statistically significant across the three outcomes we study.

rejecting the null of a similar impact across these groups, the result is consistent with the hypothesis that people most affected by the policy are more likely to support the incumbent.

Finally, columns (7)-(9) repeat the previous estimations but now use as dependent variable an indicator that takes the value of one if the person voted in the election. Estimates reveal a positive impact on the probability of turning out to vote—with point estimates similar in magnitude to previous estimates–although in this case none is statistically significant at conventional levels. All in all, these results suggest that awareness of public pharmacies and their characteristics increased consumer support for the incumbent mayor.

We combine these results with estimates of consumer savings from Section 1.5 to estimate the political returns of public pharmacies. The experiment suggests that introducing a public pharmacy increases the number of votes for the incumbent by 1,055, relative to an average of 16,105 total votes across counties in the 2012 local election. Our estimates of the effects on drug expenditure imply that the incumbent obtains 1 additional vote per \$166 of yearly consumer savings. We also consider the monthly savings of consumers who switch to public pharmacies and focus on consumers with chronic conditions. Within that population, the average individual realizes monthly savings of \$44.8. These "transfers" increased political support of the incumbent mayor by 8.1 percentage points. For reference, Manacorda et al. (2011) find that in Uruguay, a targeted monthly transfer of \$70 increased political support for the incumbent government by 11 percentage points.

1.7 Conclusion

State-owned firms compete with the private sector in a variety of markets. The costs and benefits of such competition have been difficult to evaluate empirically. In this paper, we leverage the decentralized entry of state-owned firms to a fully deregulated private market of pharmaceutical retailers. We show that the public option emerged as a low-price and low-quality option and affected the shopping behavior of local consumers, which generated market segmentation and higher prices in the private sector. Although public pharmacies created winners and losers within local markets, overall consumer savings outweighed the costs of public pharmacies.

While our study focuses on a particular form of public-private competition, it provides general lessons. First, the equilibrium effects of the public option are shaped by the nature of demand responses. In our context, the public option is less attractive to consumers with a high willingness to pay for service quality relative to drug prices. Market segmentation makes these consumers worse off due to price increases in the private sector.[38] Second, our analysis highlights the fact that public competition may be effective in reducing consumer expenditure. In industries with substantial market power in input and retail markets, retail

[38]Selection markets, like the market for health insurance, are another important context where the nature of demand responses is key for understanding the general equilibrium effects of the public option. A key feature of those settings would be whether the public option is differentially attractive to consumers with different levels of risk.

prices are set at markups over marginal costs. Whenever state-owned firms have higher bargaining power in the input market or decide not to exercise market power in the retail market, they may be able to effectively reduce consumer expenditure. Our setting indeed features these two conditions.

The political rewards of state-owned firms could be interpreted as showing that, as a whole, state-owned firms increased welfare. However, we highlight the fact that recent research shows that people may overvalue policies when they do not internalize the general equilibrium effects that affect them (Dal Bó, Dal Bó, and Eyster, 2018). Our findings are somewhat consistent with this interpretation since the majority of consumers in the market are worse off after the entry of public pharmacies due to increased private pharmacy prices.[39] These findings demonstrate the need to evaluate the market effect of policies instead of drawing conclusions about their desirability based on voting behavior.

Our analysis leaves many questions for future research. Of particular relevance is understanding the choice of quality among state-owned firms. If the quality of state-owned firms were higher, we would expect more consumers to switch to them and strengthen the equilibrium effects toward the private sector. However, changes in the quality of state-owned firms could influence their targeting properties by modifying the population that adopts them (Kleven and Kopczuk, 2011). Furthermore, it is also possible that a higher quality of state-owned firms triggers other strategic responses in the private sector. In the context of retail, these could include changes in the location, prices, or quality of private stores. Our findings thus call for attention to how the interplay between public and private firm attributes may shape equilibrium effects in the market and determine the overall and distributional impacts of state-owned firms.

[39] Recent work by Illanes and Moshary (2020) on the deregulation of retail liquor markets in Washington state also finds evidence consistent with this phenomenon.

Figures

Figure 1.1: Timing of entry of public pharmacies

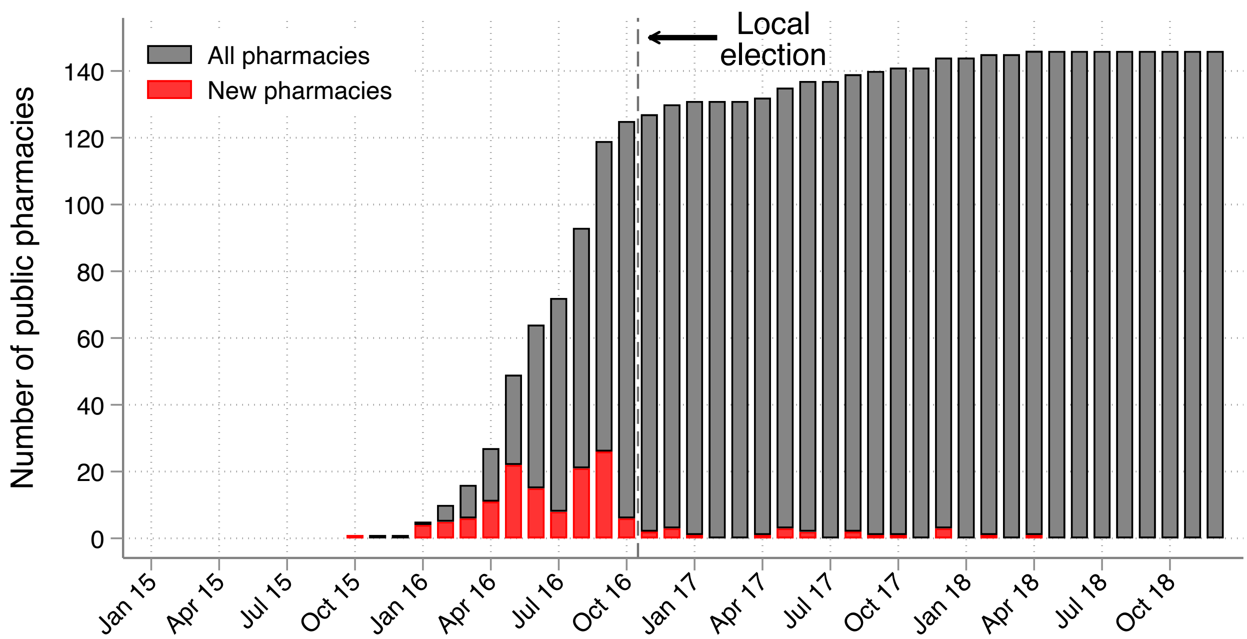

Notes: This figure shows the opening dates of public pharmacies (red bars) and the total number of public pharmacies operating (gray bars) in each month between January 2015 and December 2018. The y-axis indicates the total number of public pharmacies opened or the total number of public pharmacies operating each month during this period. The first public pharmacy opened in October 2015. The vertical dashed line in October 2016 indicates the month of the 2016 local election in which most mayors who opened public pharmacies ran for reelection.

Figure 1.2: Locations of public pharmacies in local markets

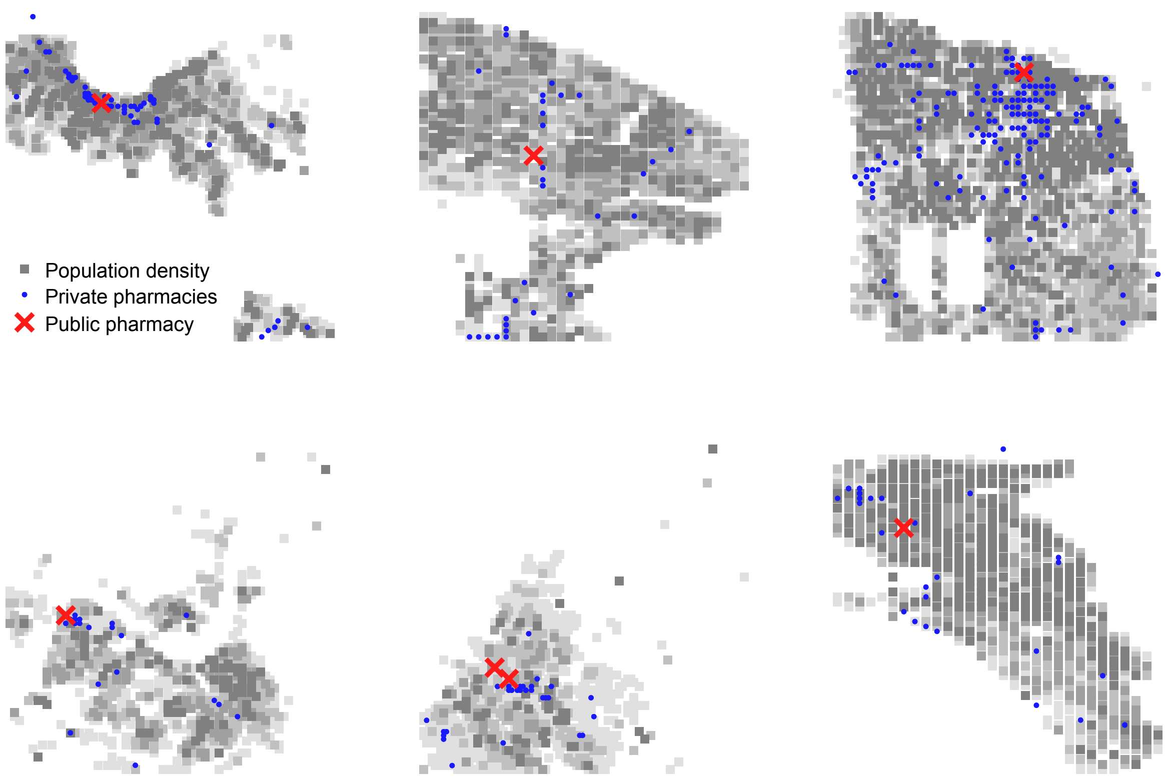

Notes: Each map represents a local market defined as a county. The maps display the exact locations of private pharmacies (blue dots), public pharmacies (red cross), and population density in cells of 111×111 meters (gray scale). White cells correspond to unpopulated (e.g., parks) or commercial areas. We categorize population density in the following five bins: $[0, 10)$, $[10, 50)$, $[50, 100)$, $[100, 150)$, and more than 150 individuals. We use the home addresses of all individuals in the country as revealed by the official Electoral Registry of 2017. The maps correspond to counties in the north, center, and south of the country: (a) Valparaiso, (b) Recoleta, (c) Santiago, (d) Valdivia, (e) Talca, and (f) Iquique.

Figure 1.3: Relative prices and attributes between private and public pharmacies

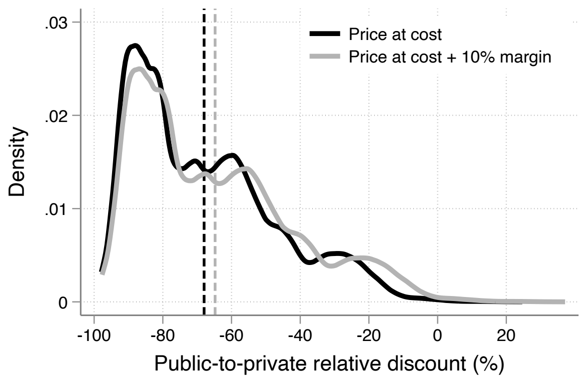

(a) Distribution of price discounts

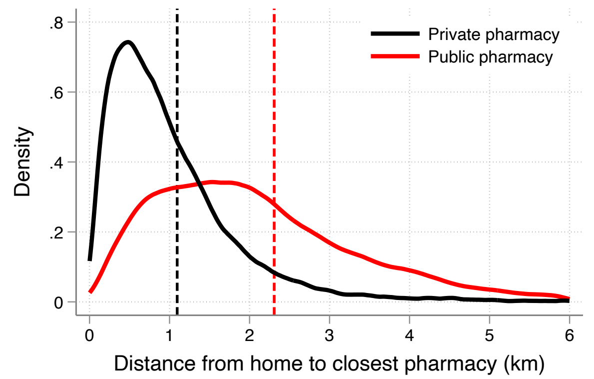

(b) Distance to pharmacies

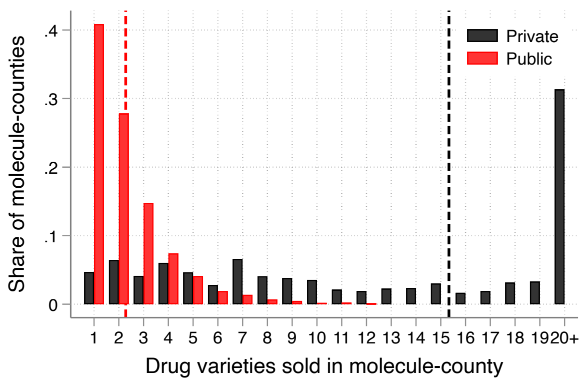

(c) Number of drug varieties

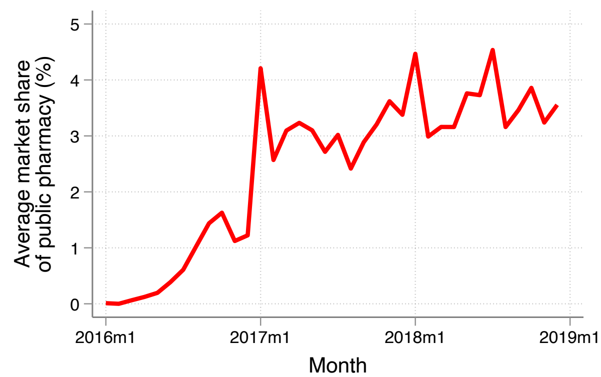

(d) Evolution of market share of public pharmacies

Notes: Panel (a) displays the distribution of proportional discounts of drugs at public pharmacies relative to private pharmacies. The plot is computed using a matched sample of the exact same drug observed in both the CENABAST (public pharmacies) and IQVIA (private pharmacies) datasets for a given county and month during 2017–2018. Because the CENABAST data only provide the cost to public pharmacies, we compute price discounts for public pharmacies pricing at cost (black) and at a margin of 10 percent over cost (gray). The dashed vertical lines indicate the mean price discount for each scenario. Panel (b) shows the density of distance from people's homes to the closest private pharmacy (black) and to the public pharmacy in counties with a public pharmacy. The dashed vertical lines indicate the respective means of both distributions. Panel (c) describes the number of drug presentations of a given molecule sold in a county over 2017–2018 for private (black) and public (red) pharmacies, whenever both private and public pharmacies sell at least one drug of the molecule. Panel (d) displays the average market share of public pharmacies across molecules and counties in each month during 2016–2018.

Figure 1.4: Impact of public pharmacies on sales and prices in private pharmacies

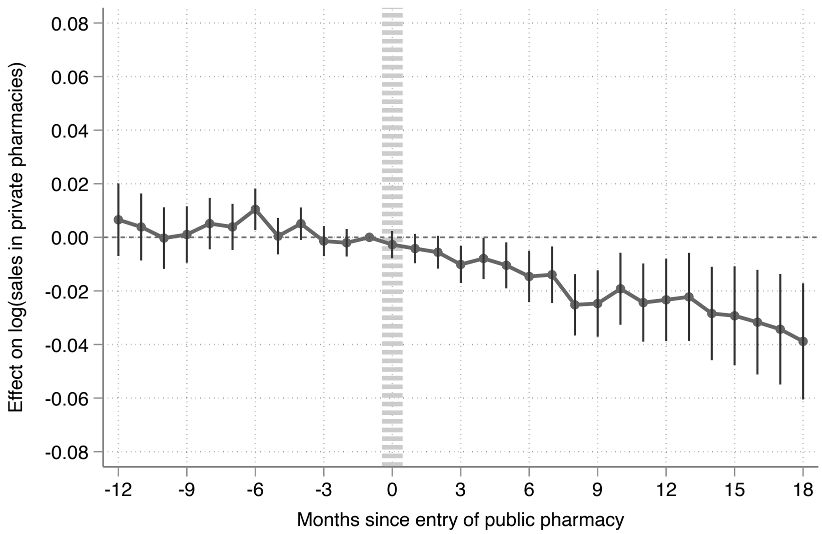

(a) Sales

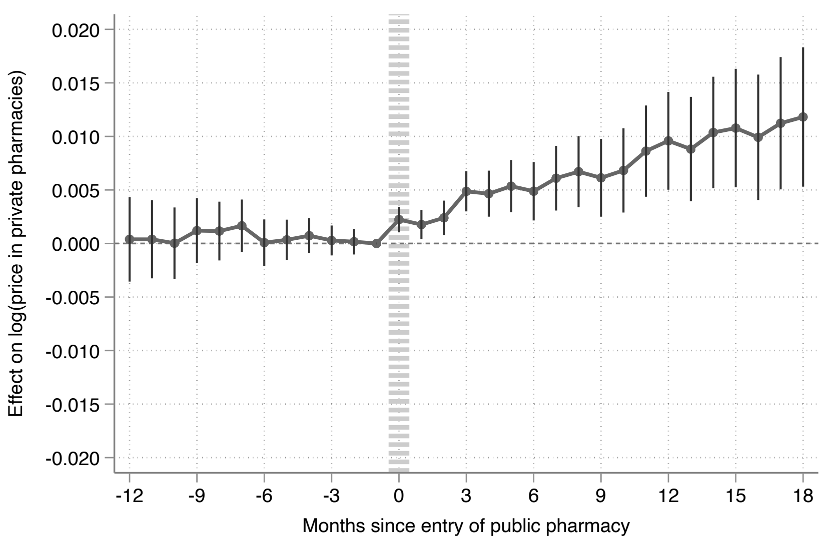

(b) Prices

Notes: These figures present event-study estimates of the impact of public pharmacies on private pharmacy sales in Panel (a), and on private pharmacy prices in Panel (b). The unit of observation is a molecule per market in a given month and we use 681,120 observations in panel (a) and 648,885 in panel (b). The empirical strategy uses panel data for the period between 2014 and 2018 and exploits the staggered entry of public pharmacies from October 2015 onward in an event-study design. In Panel (a) the dependent variable is logged sales and in Panel (b) the dependent variable is logged prices. The x-axis indicates the month with respect to the opening of the public pharmacy, i.e., 18 means 18 months after the opening, and -12 means twelve months before the opening. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95 percent confidence interval.

Tables

Table 1.1: Descriptive statistics in counties with and without public pharmacies

	(1)	(2)	(3)	(4)
	County has public pharmacy			
	Yes	No	Difference (1)–(2)	Timing of entry
Panel A: Pharmacies and hospitals				
Private pharmacies per 100,000 inhabitants	13.59	7.72	5.86***	-0.003
Log sales in private pharmacies	15.37	15.15	0.21**	-0.465
Price index in private pharmacies	931	873	59**	0.001
Hospitalizations per 100,000 inhabitants	9,440	8,126	1,313***	0.00
Deaths per 100,000 inhabitants	209	177	32***	-0.02
Panel B: Socioeconomic characteristics				
Log household income	12.97	12.61	0.36***	-0.467
Age of inhabitants	44.50	45.67	-1.18***	0.115
Average unemployment rate	0.10	0.09	0.01***	7.091
Share with public health insurance	0.83	0.89	-0.06***	1.400
Self reported health (1-7)	5.54	5.49	0.05*	1.900*
Number of doctor visits	0.32	0.30	0.02	1.359
Population (in 10,000)	9.70	1.88	7.82***	-0.425**
Panel C: Political characteristics				
Number of competitors	3.57	3.20	0.37***	0.121
Winning margin	0.19	0.17	0.02	-3.768
Vote share winner	0.54	0.53	0.01	5.951
Incumbent coalition wins	0.62	0.57	0.05	0.439
Incumbent coalition: independent	0.32	0.34	-0.03	-0.045
Incumbent coalition: left-wing	0.47	0.36	0.10*	-1.161**
Incumbent coalition: right-wing	0.22	0.29	-0.07	–
Number of counties	146	198	–	146

Notes: This table presents descriptive statistics for counties with and without a public pharmacy in 2018 in columns (1) and (2), respectively. Characteristics in panel A are own construction using data from the Public Health Institute (ISP, DEIS) and IQVIA in 2014. Socioeconomic characteristics in Panel B are own construction using data from the survey National Socioeconomic Characterization (CASEN) conducted in 2015, with the exception of "Population" data, which are publicly available on the website of the National Statistics Bureau (INE). Political characteristics in panel C are own construction using data from the Electoral Service (SERVEL). Column (3) reports the difference between columns (1) and (2) and its statistical significance. Column (4) uses the cross-section of 146 counties with public pharmacies and reports coefficients from an ordered logit using the order in which public pharmacies opened as the dependent variable—the first pharmacy has a value of 1 and the last the value of 146—and all market and political characteristics as explanatory variables. Significance level in columns (3)-(4): *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 1.2: 18-month effect on drug sales and prices in the private market

	(1) log(sales)	(2) log(price)
Panel A: Main estimates		
All sample	-0.038***	0.010***
	(0.007)	(0.002)
Panel B: Heterogeneity by chronic condition		
Molecules for chronic conditions (β_{chronic})	-0.045***	0.010***
	(0.007)	(0.002)
Molecules for non-chronic conditions ($\beta_{\text{non-chronic}}$)	-0.028***	0.011***
	(0.008)	(0.002)
p-value: $\beta_{\text{chronic}} = \beta_{\text{non-chronic}}$	0.006	0.429
Panel C: Heterogeneity by relative product variety		
High public-private variety ratio ($\beta_{\text{high variety}}$)	-0.044***	0.013***
	(0.007)	(0.002)
Low public-private variety ratio ($\beta_{\text{low variety}}$)	-0.033***	0.007***
	(0.008)	(0.002)
p-value: $\beta_{\text{high variety}} = \beta_{\text{low variety}}$	0.020	0.000
Panel D: Heterogeneity by distance to private pharmacy		
Private pharmacies are close to public pharmacy (β_{close})	-0.042***	0.008***
	(0.008)	(0.002)
Private pharmacies are far from public pharmacy (β_{far})	-0.034***	0.012***
	(0.007)	(0.002)
p-value: $\beta_{\text{close}} = \beta_{\text{far}}$	0.050	0.000
Observations	681,120	649,885
Molecule-by-month FE	Yes	Yes
Molecule-by-market FE	Yes	Yes

*N*otes: This table presents the 18-month effect of the impact of public pharmacies on private pharmacies' sales and prices. These estimates are calculated as $\bar{E}_{18} \times (\beta^{\text{jump}} + (18+1)\beta^{\text{phase in}})$, where $\bar{E}_{18}$ is the average share of population across markets with access to a public pharmacy 18 months after the first pharmacy in the local market was introduced. We estimate the on-impact effect β^{jump} and the trend break effect $\beta^{\text{phase in}}$ using an exposure difference-in-differences design that leverages the staggered introduction of public pharmacies in the panel data of molecules observed by market and month in the period 2014-2018. We report estimates of β^{jump} and $\beta^{\text{phase in}}$ in Table A.4.2. In Panel B, exposure to public pharmacies is interacted with an indicator for whether a molecule is targeted toward a chronic condition or not. In Panel C, exposure is interacted with an indicator for whether there is a high ratio of variety of products within molecule in public pharmacies relative to private pharmacies defined as above or below the median of the distribution. In Panel D, exposure is interacted with an indicator for whether private pharmacies are located "near" or "far" from public pharmacies. We use the average number of public pharmacies operating within 400 meters of private pharmacies and split the sample in two using the median of this cross-sectional market-level variable. Standard errors clustered at the molecule-by-market level are displayed in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.3: Municipal finance

	(1)	(2)	(3)	(4)	(5)	(6)
	Health services		Non-health services		All services	
	Spending	Revenue	Spending	Revenue	Spending	Revenue
Public pharmacy 18-month effect	0.041***	0.027**	-0.048	-0.032	0.015	0.010
	(0.013)	(0.013)	(0.035)	(0.033)	(0.015)	(0.014)
p-value: Spending = Revenue	0.036		0.496		0.560	
Mean of dep. var. in 2014	170.36	167.09	525.32	563.07	695.68	730.15
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Counties	321	321	322	322	321	322
Observations (county-years)	2,243	2,243	2,228	2,227	2,243	2,243

Notes: This table presents our estimates for the impact of public pharmacies on municipal finances. The health (columns 1-2) and non-health (columns 3-4) categories are mutually exclusive. Columns 5-6 correspond to "All services" provided by the county. We observe a panel of counties every year in the period 2013–2019 and exploit the staggered entry of pharmacies in a parametric event study analysis. The dependent variable is the logarithm of total spending (in U.S. dollars) per capita (2013 population) in odd columns and the logarithm of total revenue per capita in even columns. The 18-month effect is the linear combination of regression coefficients $\pi^{\text{jump}} + (1.5+1) \times \pi^{\text{phase in}}$. Table A.4.3 presents full regression results, i.e., estimates of π^{jump} and $\pi^{\text{phase in}}$. We focus on 18-month effects to compare the cost of public pharmacies with their impact on sales and prices in private pharmacies (Panel (a) of Table 1.2). Standard errors clustered at the county level are displayed in parentheses. Significance level: *** p<0.01, ** p<0.05, * p<0.1.

Table 1.4: Effect on avoidable hospitalizations associated with chronic diseases

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Avoidable hospitalizations per 100,000 inhabitants							
	Number of hospitalizations		Days of hospitalizations		Number of surgeries		Number of deaths	
Public pharmacy 18-month effect	-0.891	-1.023	-5.837	-5.061	0.116	0.075	0.092	0.123
	(0.788)	(0.826)	(7.815)	(8.527)	(0.175)	(0.195)	(0.084)	(0.093)
Health insurance	All	Public	All	Public	All	Public	All	Public
Mean of dep. var. in 2014	17.93	19.18	158.1	172.5	1.724	1.907	0.736	0.828
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Counties	344	344	344	344	344	344	344	344
Observations (county-month-years)	28,320	28,320	28,320	28,320	28,320	28,320	28,320	28,320

Notes: This table presents our estimates for the impact of public pharmacies on avoidable health outcomes. The outcomes of interest are the number of hospitalizations (columns 1-2), days of hospitalizations (3-4), number of surgeries (columns 5-6), and number of deaths (columns 7-8). For each outcome, the first column uses the count of the outcome per 100,000 inhabitants in a county regardless of individual health insurance, and the second column restricts that count to individuals with publicly provided insurance (FONASA). We observe a panel of counties every month in the period 2013–2019 and exploit the staggered entry of pharmacies in a parametric event study analysis. The 18-month effect is the linear combination of regression coefficients $\pi^{\text{jump}} + (18+1) \times \pi^{\text{phase in}}$. Table A.4.4 presents full regression results, i.e., estimates of π^{jump} and $\pi^{\text{phase in}}$. We focus on 18-month effects to use the same horizon of effects as in the previous estimates in the paper. We report the mean of the dependent variable for 2014 among counties that ever introduce a public pharmacy, the year before most public pharmacies entered the market. Standard errors clustered at the county level are displayed in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 1.5: Experimental results for political outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Voted incumbent mayor			Voted incumbent party			Voted in the election		
Treatment	0.057	0.075*		0.064	0.056		0.066	0.052	
	(0.045)	(0.039)		(0.040)	(0.035)		(0.046)	(0.044)	
Treatment × chronic (β_C)			0.080			0.081*			0.040
			(0.051)			(0.044)			(0.055)
Treatment × non-chronic (β_{NC})			0.067			0.020			0.068
			(0.065)			(0.058)			(0.073)
Dependent variable at baseline		0.366***	0.367***		0.348***	0.350***		0.418***	0.416***
		(0.051)	(0.051)		(0.048)	(0.048)		(0.052)	(0.052)
Lee bounds	[0.033, 0.182***]			[0.048, 0.170***]			[0.014, 0.159**]		
p-value for H_0: $\beta_C = \beta_{NC}$	-	-	0.883	-	-	0.408	-	-	0.763
Mean for control group	0.281	0.277	0.277	0.263	0.255	0.255	0.541	0.524	0.524
Observations	398	368	368	475	435	435	475	435	435
R-squared	0.004	0.515	0.515	0.005	0.488	0.488	0.004	0.641	0.641
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
County FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes

Notes: This table presents our estimates of the political impact of public pharmacies using data from the field experiment described in Section 1.6 The unit of observation is an individual who buys pharmaceuticals at private pharmacies in the capital city of Santiago. The treatment is information about public pharmacies delivered in the form of a flyer by enumerators after completing the baseline survey in October 2016, before the local election. All dependent variables were measured in follow-up surveys conducted in December 2016, after the local election. We present cross-sectional results using three specifications, one without controls (columns 1, 4, and 7), one with controls (columns 2, 5, and 8), and one with controls and interacting the treatment with an indicator for individuals with a chronic condition (columns 3, 6, and 9). The set of control variables includes age and indicators for chronic condition, having completed high school education, female, and public insurance. Reported Lee bounds are computed using only the treatment indicator as covariate. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Chapter 2

Managers and Public Hospital Performance

2.1 Introduction

Global government spending on publicly provided goods and services more than doubled between 1980 and 2019, and accounts for approximately 30% of world GDP (Gethin, 2022). Given the scale and scope of this spending, enhancing state efficiency is an important channel for increasing total productivity. Government policies intended to strengthen state efficiency often focus on improving the job performance of public sector managers, who directly supervise the delivery of goods and services provided by the state and are uniquely situated to advance the goals of the state (Pollitt and Bouckaert, 2017). However, research on whether and how public sector managers can improve their organization's performance is limited (Bertrand, Burgess, Chawla, and Xu, 2020; Fenizia, 2022). Empirical progress is difficult because of two methodological challenges. First, quasi-experimental variation in state personnel selection processes is rare. Second, it is challenging to study managerial effectiveness due to the lack of objective and verifiable performance outcomes in the public sector (Besley, Burgess, Khan, and Xu, 2022).

In this paper, we study the impacts of public sector managers in a unique setting that allows us to overcome these methodological challenges and to shed light on key drivers of public sector productivity. Specifically, we analyze a novel policy in Chile that reformed the selection process for senior executives in the public sector. Although this reform applied broadly across all public agencies and departments, we focus on the top managers (CEOs) of public hospitals, which allows us to observe objective and relevant short-term outcomes to assess their managerial performance. Focusing on public sector hospitals also allows us to study a setting in which government expenditures are large and growing, outcomes for patients are high-stakes, and disadvantaged communities are particularly likely to interact with the public sector.[1]

[1]Healthcare represents almost 20% of government expenditures in the average OECD country, and between

The policy reform that we study was enacted by the Chilean Congress in 2003, and resulted in the introduction of a new personnel recruitment system designed to attract talent to top management positions in the public sector. The reform had two main components. First, it replaced an opaque and discretionary selection process with a public, competitive, and transparent selection system for senior executive positions. Second, it included performance pay incentives and base wage increases to narrow compensation differentials relative to similar positions in the private sector. The reform affected top-level positions in public agencies and was gradually implemented across all ministries and other public organizations. In 2004, eight managers in senior executive positions were hired using the new selection system; by 2019, the new system was used to appoint more than 3,400 senior executives in 1,400 positions.

To study the impacts of public sector CEOs on hospital performance, we build a novel and comprehensive dataset with information on the identity, tenure, educational background, cognitive skills, and demographic characteristics of CEOs in all public hospitals between 2001 and 2019. We complement these data with restricted-use employer-employee matched data from the Ministry of Health for all public hospitals, which allows us to observe individual-level wages and employment. To measure hospital performance, we use data on nationwide individual-level inpatient discharges for all public hospitals, and individual-level death records in the whole country. We complement these data with hospital inputs and procedures. The data thus provide an extraordinarily rich window into hospital mortality, procedures and inputs, patient characteristics, and the characteristics and tenure of CEOs in every public hospital in Chile.

We begin by documenting that CEO managerial talent matters for hospital performance. We follow the literature in using hospital mortality rates as our key measure of outcome-based productivity (e.g., Bloom, Propper, Seiler, and Van Reenen, 2015b; Doyle, Graves, Gruber, and Kleiner, 2015; Chandra, Finkelstein, Sacarny, and Syverson, 2016; Doyle, Graves, and Gruber, 2019) and show that individual CEOs can account for a significant amount of variation in mortality rates across hospitals. We follow the approach pioneered by Bertrand and Schoar (2003), and gradually include hospital and CEO fixed effects in a regression on hospital mortality. Inclusion of CEO fixed effects increases the R^2 by a similar magnitude to that in Bertrand and Schoar (2003) for CEOs in publicly traded US firms and in Fenizia (2022) for managers in the administrative public sector in Italy.[2]

We next study the effects of the selection reform when adopted by public hospitals for appointing their CEOs. We present three main findings. First, we find that the reform significantly improved hospital performance. We exploit the gradual adoption of the new

2000 and 2019 healthcare costs increased by 15% as a share of GDP on average in OECD countries and have boomed in the post-pandemic world. In the OECD, public hospitals contribute, on average 72% of the total supply of medical beds (see Appendix Figure B.6.1).

[2]In Appendix B.2, we further exploit the rotation of CEOs across hospitals and estimate a model with CEO and hospital fixed effects. This procedure allows us to compute measures of managerial talent (CEO fixed effects) and to decompose the variance of mortality and quantify the contribution of fixed hospital characteristics and CEO talent.

selection system for public hospital CEOs to estimate its causal effects on hospital performance using a staggered difference-in-differences research design. We show that the reform decreased death rates by approximately 8% in public hospitals in the 3 years following adoption of the new system. These effects are similar to the mortality impacts of other policies studied in the literature, such as increasing patient expenditures by 10% (Doyle et al., 2015) or improved clinical practices in VA hospitals in the United States (Chan, Card, and Taylor, 2022).

Our empirical analysis is subject to two econometric concerns. The first relates to the identification assumption underlying the staggered difference-in-differences model, namely that hospitals that adopt the selection reform are not on trends that differ from those that have not adopted it. We justify this assumption in several ways. First, we show that in the years before the first hospital adopted the reform, the growth of an exhaustive set of variables—including hospital outcomes, patient characteristics, and political variables—do not differ between hospitals that do and do not adopt the new selection process. Second, using an event study design, we show graphically that hospitals are not on different trends in mortality rates before the adoption of the policy. The lack of pre-trends eases concern regarding an Ashenfelter-style dip, which is a natural threat in settings in which management changes can respond to a decline in performance. Third, we show that the dynamic effects of the reform gradually grow during the early quarters post-adoption and flatten after, which is the expected trajectory if new managers are to have an impact on performance. Additionally, we provide event study evidence showing that CEO transitions in hospitals that have not adopted the selection reform have zero impact on hospital performance, which rules out mechanical effects of the reform due to CEO turnover.

The second concern is that our estimates might reflect changes in patient composition rather than improved CEO performance: Perhaps after the reform, managers admitted healthier inpatients, or perhaps patients self-selected into hospitals with improving performance. We provide several pieces of evidence to address this concern. First, we underscore that the Chilean public health system is particularly well suited for this study because there is minimal scope for patient selection. Within the public health network, patients cannot choose their hospital provider and are referred from primary care centers to public hospitals following strict guidelines. By the same token, hospitals cannot select patients based on their characteristics (Ley 19,937; Decreto 38). Consistent with the setting's features, we do not find any evidence of impacts of reform adoption on the hospital-level risk scores predicted based on patients' demographics and diagnoses. Second, in our baseline estimates we use an exhaustive set of case mix controls that include detailed information on patient demographics and diagnoses. We also provide estimates for risk-adjusted mortality rates following the prediction procedure used by the Centers for Medicare and Medicaid Services (CMS) in the United States. Third, to deal with further concerns regarding selection on unobservables, we examine the effect of the reform on deaths outside treated hospitals. To the extent that patients rejected by a given hospital die, they would show up in the statistics of other hospitals or be recorded as home deaths. We find evidence of zero effects in neighboring hospitals or aggregate home deaths at the municipality level. Finally, we also find that the reform had similar effects when we focus exclusively on "locked-in" patients who cannot access healthcare in the private sector, or when we restrict the analysis to the set of patients that followed the referrals mandated by the health system.

Consistent with the effects on hospital mortality, we find that the reform induced a more efficient use of medical resources and improved personnel practices. Operating room utilization increased by 25% three years after the reform adoption, a large effect, but not large enough to close the gap between the average efficiency in high-complexity hospitals in our setting and the average efficiency in the National Health System (NHS) of the United Kingdom. In line with this finding, we also show that surgical procedures increased by a similar magnitude; and that most of this increase was led by surgeries in diagnoses amenable to death prevention through surgery, such as cardiac arrhythmia, renal failure, and metastatic cancer. Despite finding null effects on staff compensation, the reform significantly reduced the turnover of doctors. These results are consistent with findings by Bloom et al. (2015b), who document a positive correlation between improved management practices, hospital performance, lower waiting times, and reduced staff turnover. The results are also in line with recent research in personnel economics documenting that better-managed firms retain workers with higher human capital (Bender, Bloom, Card, Van Reenen, and Wolter, 2018).

The second finding is that the reform's new financial incentives for CEOs—namely, performance pay and higher base wages—did not play a role in driving the mortality effects. We start by ruling out the possibility that performance pay affected managerial performance. We show that incentives were poorly designed and were not binding, a feature that was true across all positions appointed using the reform's selection system and not specific only to public hospitals. We also document that the policy increased CEO wages by approximately a third relative to pre-reform wages. To examine whether these higher wages had an effect on managerial performance, we leverage an amendment to the reform that increased the pay for managers who are doctors *and* who were appointed after November 2016. We find that the amendment significantly increased wages for treated managers but had no discernible effects on their performance. We interpret these results as evidence that efficiency wages do not drive the impact of the reform on hospital mortality.

Our third finding is that the reform dramatically increased the share of CEOs with management education, which we show is the main predictor of the efficacy of the reform on reducing hospital mortality. We first document that the reform replaced doctors who worked as CEOs ("doctor CEOs") with new CEOs with undergraduate degrees in management-related majors.[3] Before the reform, 93% of CEOs were doctors. After the reform, the share of CEOs with undergraduate degrees in management increased, on average by 21 percentage points, and the share of doctor CEOs decreased by a similar magnitude. Interestingly, the reform had a large and negative effect only on doctor CEOs without management training while it substantially increased the share of doctor CEOs with managerial qualifications. We present suggestive evidence that through this channel, the reform also had an effect by incentivizing

[3]Management-related majors include public administration, business and economics, accounting, and engineering.

doctors who wanted to apply for a CEO position to further invest in management studies. Overall, we find that, soon after it was adopted, the new selection process boosted the percentage of CEOs with management training by over 40 percentage points. We then show that the reform increased CEO managerial talent—as proxied by their estimated CEO fixed effect—, but had no significant effect on cognitive abilities—as measured by standardized college entrance exams scores. We also find that new managers are approximately 2 years younger and that the reform had no effect on the likelihood that the CEO is female.

Motivated by these findings, we next examine whether correcting the mismatch between CEOs' skills and the skills demanded by the job enhanced the organization's performance. We interact the reform with CEO management training and find that the beneficial effects on hospital mortality are primarily driven by CEOs with managerial qualifications. To further explore this result, we estimate stacked event studies leveraging all CEO transitions from CEOs without management training to CEOs with management training.[4] We find that these transitions consistently produce a significant and large decline in mortality rates, while transitions between CEOs with no management training have zero effect on hospital performance. While our results are consistent with the hypothesis that management training improves the performance of hospital CEOs, we cannot rule out that the effect is explained by differential selection (i.e., better managers are more likely to study management).

The norm before the introduction of the policy was that CEO positions in public hospitals were overwhelmingly reserved for doctors; in fact, doctors comprised 98% of all public hospital CEOs.[5] This pattern, where top executives rise up from the lower ranks of their profession, remains ubiquitous in public sector organizations such as police departments, school districts, and universities. Our results suggest that management skills are transferable across organizations and, importantly, that management training reaps benefits for public sector organizations even when managers acquire that training later in life. A natural policy implication is that public sector organizations may wish to emphasize management education when recruiting executives, even if those candidates rise up from the lower ranks of their respective professions.

This paper contributes to several strands of related literature. We contribute to the literature on the impacts of discretionary appointments (Myerson, 2015; Padró i Miquel, Qian, and Yao, 2018; Xu, 2018; Colonnelli, Prem, and Teso, 2020; Voth and Xu, 2022) and civil service recruitment in the public sector (Dal Bó, Finan, and Rossi, 2013; Muñoz and Prem, 2022; Moreira and Pérez, 2022; Dahis, Schiavon, and Scot, 2022). To the best of our knowledge, this is the first paper to evaluate the impact of state personnel selection processes on health outcomes. More broadly, our research also contributes novel evidence

[4] To avoid a mechanical impact of the reform, we exclude transitions induced by the reform adoption.

[5] Based on qualitative interviews, we posit that the norm was upheld by a shared belief across doctors that individuals with no medical training should be largely excluded from from CEO positions and should only be permitted to advance to middle management. According to them, there is an insurmountable information asymmetry between the medical staff and non-doctors, which renders non-doctors unqualified to make the final decision about how to run a hospital. Furthermore, biased beliefs about the importance of management may have previously discouraged doctors from investing in management training.

on bureaucratic effectiveness and its impact on development (see Besley et al. (2022) for a review).

Our work also adds to a nascent literature studying the impact of top managers (Choudhury, Khanna, and Makridis, 2019; Limodio, 2021; Fenizia, 2022; Best, Hjort, and Szakonyi, 2022) and management practices (Bloom, Lemos, Sadun, and Van Reenen, 2015a; Tsai, Jha, Gawande, Huckman, Bloom, and Sadun, 2015; Rasul and Rogger, 2018; Azulai, Rasul, Rogger, and Williams, 2020; Bloom, Lemos, Sadun, and Van Reenen, 2020) on organization-level performance in public sector organizations.[6] Closest to our work, Janke, Propper, and Sadun (2020) study the impact of CEOs in NHS hospitals in the UK and document little evidence of CEOs' impact on different dimensions of hospital production. In contrast to our setting, CEO recruitment in NHS hospitals does not have a strict selection criteria and is delegated to local boards.[7] By studying a selection reform that dramatically shifted CEOs' managerial qualifications, we also provide novel evidence on the impact of correcting top managers' skill mismatches in public sector organizations.

This paper also contributes to the literature on management styles (e.g., Hambrick and Mason, 1984; Bertrand and Schoar, 2003). In contrast to recent work by Acemoglu, He, and le Maire (2022) that focuses on private sector firms, we find null effects of appointing CEOs with management studies on the wages of public employees. Instead, we document that performance was improved in hospitals where the new CEOs had management training.[8] Moreover, our study of the impact of CEOs on personnel turnover adds to existing work on management and employee attrition (Hoffman and Tadelis, 2021), and underscores its importance for high-skilled workers (Bender et al., 2018; Gosnell, List, and Metcalfe, 2020). Finally, our work complements previous research studying how to improve operating room efficiency (e.g., He, Dexter, Macario, and Zenios, 2012) and the efficacy of alternative policies to improving hospital performance (e.g., Propper and Van Reenen, 2010; Gaynor, Moreno-Serra, and Propper, 2013; Bloom et al., 2015b; Doyle et al., 2019; Chan et al., 2022; Duggan, Gupta, Jackson, and Templeton, 2022).

The rest of the paper proceeds as follows. Section 2.2 briefly describes the setting and data, provides descriptive evidence of CEO impact on hospital performance, and introduces the recruitment reform. Section 2.3 presents the main effects of the reform on hospital mortality, and discusses the validity of the results and potential mechanisms. Section 2.4 examines the effects of the financial incentives included in the reform. Section 2.5 examines

[6] Management's effects on organizational performance can operate through the manager herself, organizational-level management practices or a combination of both (see Metcalfe, Sollaci, and Syverson (2022) for discussion).

[7] Interestingly, in the NHS only a quarter of CEOs hold postgraduate management education, which is the same to the average in our setting before the selection reform adoption within hospitals that adopted it. The reform increased the share of CEOs with postgraduate management qualifications to almost 60% the quarter after adoption (and up to 66% if we also consider undergraduate management degrees).

[8] In terms of methodology, our paper departs from previous field experimental literature focusing on management practices Bloom et al. (2015b) and connects with the literature leveraging CEO turnover such as Bennedsen, Pérez-González, and Wolfenzon (2020) and Acemoglu et al. (2022).

the recruitment effect of the reform on managers' characteristics. Section 2.6 concludes and the Appendix provides additional results.

2.2 Setting, Data, and Descriptive Evidence

The Healthcare System in Chile

Chile's healthcare system comprises public and private health providers and public and private insurers. Public insurance is funded by general taxation and payroll taxes on enrolled employees. Individuals can opt-out and use their health contributions to buy private insurance.[9] Individuals without the ability to pay can freely access the public system, which results in nearly universal health coverage.

Approximately 78% of the population is under public health coverage, 15% have private insurance, and the remainder are covered under special regimes exclusive to the police and armed forces.[10] The ability of individuals to freely use their health contribution to buy private coverage has induced stark market segmentation, because private insurers are able to charge differentiated premiums and select healthier and more affluent customers. While this has led to massive sorting across the private and public health sectors,[11] there is little scope for selection within the public health sector; the reason is that individuals with public coverage cannot choose their health provider within the public network.[12] Individuals need to register in the healthcare center that provides primary care in their local area and patients who need specialized attention are referred to specialty clinics or a hospital. Referrals follow strict guidelines, mostly based on the geographic location of the patient's primary care center (Ley 19,937). In Appendix B.1, we describe the referral process and empirically show the lack of selection within the public network. Patients can also be admitted directly to public hospitals in emergency cases.

Public healthcare providers are organized geographically under 29 "Health Services," the administrative units within which the referral and counter referral system is organized. These are decentralized organizations subject to oversight by the Ministry of Health, and

[9] The healthcare system in Germany features an analogous mechanism. Upon meeting certain conditions, individuals can use their health contribution to buy private insurance (known as PKV) and opt-out of the public health insurance system (known as GKV).

[10] For comparison, private compulsory health insurance spending explains approximately 10% of health expenditures, similar to Germany and France. In 28 out of 35 OECD economies, however, it comprises less than 5% of health expenditures (OECD, 2022b).

[11] Almost 70% of people in the top 10% of the income distribution have private coverage, while only 4% in the bottom 50% buy private coverage (CASEN, 2017).

[12] While private insurers may provide coverage in public hospitals, this is rarely seen in the data. The reason is that individuals under private insurance are already self-selected into the private health sector and have little incentive to choose public healthcare providers. In the universe of admissions, 96% of patients at public hospitals have public insurance. Under public coverage, individuals can choose private health centers, although it is more expensive than public hospitals. Around 25% of inpatients at private hospitals have this coverage.

each is responsible for the articulation, management, and development of public healthcare establishments in municipalities in their territory. This includes primary, secondary, and tertiary public healthcare and other private establishments that maintain agreements with the respective Health Service. Appendix Figure B.6.2 shows the geographic distribution of the 29 Health Services and the municipalities within their scope.

The head of the Health Service is also the immediate superior of CEOs of public hospitals within their territory. CEOs are in charge of the management, organization, and administration of their hospital, and their duties include, among others, (i) the administration of personnel, (ii) the allocation of hospital inputs and human resources, (iii) the management of financial resources and proposing the annual budget, (iv) infrastructure and technological equipment resources decisions, and (v) integration of the hospital into the health network and with the community, among others.

Data Sources

For this paper, we build a novel dataset that identifies the CEO in every public hospital in the country, spanning every month between January 2001 and December 2019. Because these data were not available in a systematic way, we filed nearly 1,000 Freedom of Information Act (FOIA) requests to local hospitals and health authorities—who, in many cases, had to collect archived data. We complement these data with background and performance records. For background characteristics, we collected date of birth, gender, test scores, and educational attainment. We gather this information from several sources, including a national registry of all medical personnel in the country, CVs requested by the Civil Service, LinkedIn profiles, articles from local newspapers, and information provided by universities, among others. Finally, via a series of FOIA requests to the Civil Service, we also have access to pay-for-performance agreements and job performance scores for CEOs appointed under the new selection system.

We also access restricted-use administrative records that cover the universe of employees in all public hospitals between 2014 and 2019. The data are collected by the Ministry of Health and unified in a single registry for HR purposes, the "Human Resources Information System." Data include detailed payroll information and wages at the monthly level. Among other characteristics, we observe the establishment, the person's job (and, in the case of doctors, their specialty), number of hours worked, date of birth and gender, and a detailed wage breakdown.

In terms of hospital performance, we use detailed administrative data collected by the Ministry of Health (DEIS, 2019). We access individual-level inpatient events that end in a discharge or death in all public hospitals in Chile between 2001 and 2019, which encompasses almost 29 million events. Data include the diagnosis according to the 10th edition of the International Classification of Diseases (hereafter, ICD-10 codes); the type of admission (e.g., emergency case or referral); the date of discharge or the date of death in case the individual died in the hospital; and a set of individual characteristics, including date of birth, gender, county of residence, and type of health insurance. We link the data at the individual level with country-wide death records processed by the Vital Records Office between 2001 and 2018, which cover more than 1.5 million deaths. We observe the date, cause, and place of death. Finally, we also collect a host of inputs and procedures at the hospital level, such as the number of medical beds, the number of surgeries, and hours of operating room use, among others. These data come from the REMs ("Resúmenes Estadísticos Mensuales") collected by the Ministry of Health, starting in 2009. We complement the data with a set of characteristics that describe the hospital, such as hospital size, whether it is public or not, and location, among others.

Finally, to compute the timing of the policy, we use data on all recruitment processes conducted by the Civil Service, which are publicly available on their website. The information includes the recruited individual's identity, the appointment date, and the Ministry in which the agency and the position are based.

Hospital mortality and CEO performance

Our main outcome of managerial performance is hospital mortality, which the literature uses extensively to measure outcome-based hospital quality in different settings (e.g., Geweke, Gowrisankaran, and Town, 2003; Gaynor et al., 2013; Bloom et al., 2015b; Doyle et al., 2015; Chandra et al., 2016; Hull, 2020; Gupta, 2021; Chan et al., 2022). A critical concern, however, is that hospital death rates might reflect shifts in the observed and unobserved characteristics of patients, potentially biasing the results of the analysis. The Chilean public health setting is well suited for this analysis because the selection of patients is limited by the institutional design. Public hospitals receive patients following strict referral guidelines based on the patient's county of residence, age, and diagnosis. Also, hospitals cannot reject patients or discretionally counter-refer them to other hospitals, and must abide by the protocols.[13] We provide further details in Appendix B.1.

We start by studying the extent to which variation in hospital quality can be explained by CEO managerial talent. Specifically, we compare the adjusted R^2 estimated from a regression of the logarithm of death rates on different sets of explanatory variables including CEO and hospital fixed-effects. We report the results in Table 2.1. Column (2) excludes hospital and CEO effects, column (3) adds hospital effects, and column (4) includes CEO effects. The adjusted R^2 increases from 0.42 in column (2) to 0.67 in column (3), which implies that hospital effects account for substantial variation in the outcome. It further increases to 0.76 in column (4) after inclusion of CEO fixed effects—an increase of similar magnitude to that reported by Bertrand and Schoar (2003) and Fenizia (2022).[14] Formally, an F-test strongly rejects the null hypothesis that all the CEO effects are zero (p-value=0.00).

[13] It might be contested that CEOs could change the referral protocols in their hospitals to avoid sicker patients. However, the referral and counter-referral system for each hospital is set and revised by the Health Service where the hospital is based and is approved by Subsecretaría de Redes Asistenciales.

[14] This finding stands in contrast to Janke et al. (2020) who—in the context of public hospitals in the English National Health Service (NHS)—document lack of CEO effects in hospital production, despite substantial and persistent differences in their pay.

In light of research casting doubt on this type of approach (Fee, Hadlock, and Pierce, 2013), in Appendix B.2 we also assess the relative importance of hospitals and managers estimating a two-way fixed effects model, which allows us to perform a variance decomposition analysis.[15] We identify the model using the rotation of CEOs across hospitals, in the same spirit as the rotation of workers identifies worker and firm fixed effects in Card, Heining, and Kline (2013).[16] Using bias-corrected measures of the variance components (Andrews, Gill, Schank, and Upward, 2008), we find that CEO fixed effects explain 26% of the variation in mortality, a magnitude similar to that of the permanent component of productivity associated with different hospitals (36%). We also find that the (bias-corrected) covariance between CEO and hospital fixed effects is negative, suggesting the best managers work at the least productive hospitals.

The Recruitment Reform

In 2003, a political scandal exposed illegal payments to top government officials. In response to and as a product of broad political consensus, Congress enacted Law Nº 19,882, which created a new framework to regulate the public sector's personnel policy (Ley 20,955). Under this new framework, the law created the Senior Executive Service System with the aim "to provide government institutions—through public and transparent competitions—with executives with proven management and leadership capacity to execute effectively and efficiently the public policies defined by the authority."[17]

The reform has two main components. First, it changes the recruitment process for top managers in government agencies. Before the reform, most senior executive positions were discretionary appointments by the superior officer. After the reform, top managers are selected through public, competitive, and transparent competitions.

The job announcement for a top management position starts with the position's being posted online on the Civil Service's website and in a newspaper with national circulation. In some cases, the Civil Service may also hire headhunters to widen the pool of applicants. Applicants must have a professional degree and, depending on the position, other competencies may be desired. After the job posting closes, the Civil Services sends the set of eligible applicants to a third-party human resources firm that evaluates each individual's job trajectory according to the job profile. They also evaluate candidates' motivation and overall competencies. The consultant assigns every applicant a grade based on an objective rubric and provides a short list to the Civil Service. In the next phase, a committee formed

[15] This model also provides us with estimates of CEO fixed effects, which are a useful measure of managerial talent we use throughout the paper.

[16] Models with additive hospital and manager components may raise some concerns. One may worry, for instance, that managers are assigned to hospitals on the basis of unobserved factors that determine their comparative advantage. It could also be that manager rotation is correlated with hospital-specific trends. Following Card et al. (2013) we empirically test these concerns and find no evidence to support them.

[17] According to the Civil Service's statement of the purpose of the reform, available at https://www.serviciocivil.cl/sistema-de-alta-direccion-publica-2/.

by representatives of the Civil Service and the Ministry in which the position is based interviews the remaining candidates and selects a short list of three individuals based on objective criteria. In the last step, the superior officer selects the winning candidate from the final roster with discretionary authority. Appendix Figure B.6.3 provide a visual illustration of the recruitment process.

The reform also increased CEO pay by providing higher base wages and performance incentives. The size of the wage increase varies across positions and is paid as a monthly bonus.[18] In the case of public hospital CEOs, we document that the reform increased the position's pay by 33% (see Appendix B.5 for further details). The financial package also includes a performance pay component, under which the yearly wage is penalized if the manager does not meet certain performance thresholds. We provide more details of the performance pay schedule and performance scores in Section 2.4.

Adoption of the recruitment process occurred gradually across public agencies over time. The law mandated that between 2004 and 2010, the Ministry of Finance had to determine a minimum of 100 top executive positions to adopt the new recruitment system. Panel A in Figure 2.1 depicts the number of positions in public agencies that adopted the recruitment reform between 2004 and 2019. All new top management positions created after the law was enacted must select their top manager using the new selection system. For existing positions, once they are subject to the new recruitment system, all future managers must be hired by the same process (i.e., treatment is an absorbing state). Panel B in Figure 2.1 shows the number of recruitment processes conducted by the Civil Service in a given year. The spikes we observe in 2011, 2015, and 2019 are evidence of substantial turnover in senior executive positions after a new government is in place.

Each adoption is costly, and therefore the Ministry of Finance has to approve it. The reason is twofold. First, the Civil Service has constrained capacity and can oversee only a limited number of processes without increasing its personnel. Second, adopting the recruitment process for a position implies higher wages and the costs of running the process (which include, among others, hiring a certified human resources firm to lead part of the selection process). Since adopting the reform triggers the new selection process for all future managers, each adoption implies a permanent expense.

In the case of public hospitals, adoption is mainly driven by their size and complexity: high, medium, and low, which is defined by the number of beds and the number of procedures they offer. Note that when the Ministry of Finance approves the recruitment process for a given position, it only takes effects after a manager transition. Therefore, the timing is also explained by transitions of incumbent managers. In Appendix Figure B.6.4 we plot a histogram of the adoption of the recruitment policy in public hospitals between 2005 and 2019. The first time a public hospital adopted the selection system was during the fourth

[18]Two limits cap the extra bonus. First, it cannot be larger than 100% of the base wage (which in the public sector is substantially lower than the total remuneration due, for example, to other tenure- and sector-specific bonuses). Second, the total wage cannot be higher than that of the Under Secretary of the Ministry in which the position is based.

quarter of 2005, after which other hospitals adopted it gradually over time. In total, 88 out of 188 hospitals adopt the selection reform during the time window of the study.

Sample and Descriptive Statistics

We use records on the universe of public hospitals overseen by the network of Health Services and aggregate the data at hospital-by-quarter level for the analysis. Aggregating the data for each hospital at quarter level is useful to avoid observations with too few discharges and to reduce volatility in the data, but results are robust to alternative aggregations. We start by constructing death indicators at patient level following a hospitalization event. We merge inpatient and death records, regardless of whether death occurred in the hospital or at another location. It is important to observe the effect on deaths outside the hospital in the analysis, because in-hospital deaths could miss patients who die shortly after discharge (Gaynor et al., 2013). We construct the hospital mortality rate as the share of inpatients who either died in the hospital or died outside the hospital 28 days after admission. Since we can follow individuals over time, we also compute death rates for different time horizons after discharge, which will be useful for performing benchmark comparisons of our results with the literature.

Our final sample consists of 188 public hospitals—of which 88 adopted the recruitment reform at some point between 2004 and 2019—for a total of 13,988 observations of hospitals-by-quarter. Table 2.2 presents descriptive statistics. In our sample, 33% and 15% of hospitals are classified as high- and medium-complexity hospitals, respectively. The average hospital in our sample discharges 1,491 patients per quarter, while the median hospital discharges 587 patients. On average, 59% of these discharges correspond to female inpatients and 36% to inpatients younger than 29 years. Importantly, 96% of patients discharged from public hospitals have public insurance. Regarding hospital outcomes, the average hospital experiences 38 deaths per quarter, with a corresponding in-hospital death rate of 2.46%. Relatedly, the 28 days after admission death rate—which considers both in- and out-of hospital deaths—is larger and corresponds to 4.21%. Finally, regarding emergency room admissions, the average death rate for ER patients is 3% when considering all diagnoses and 12.2% when considering only ER admissions with an acute myocardial infarction diagnosis.

2.3 The Reform's Impact on Hospital Performance

Research Design: Reform Adoption in Public Hospitals

Public hospitals that adopted the selection reform differ systematically from those that did not. However, the growth of a set of variables before the reform was passed is not clearly correlated with whether the hospital adopted the reform. This feature allow us to use the adoption as a plausible source of exogenous variation to estimate the impact of the reform on performance outcomes.

We compare the characteristics of hospitals that adopted the selection reform at some point (ever-treated) to the characteristics of hospitals that never adopted it (never-treated). For ever-treated, we consider a window of six quarters before adoption. Column (1) in Table 2.3 shows the average at never-treated hospitals for a set of variables related to patient demographics, hospital outcomes, and political outcomes at the hospital's municipality. Column (2) presents the OLS estimate on a dummy that takes value 1 for ever-treated hospitals and 0 otherwise. We find that on average, adopters have higher death rates and served patients who are slightly younger and less likely to use public health insurance; they are also located in municipalities that exhibited more support for right-wing politicians in the 2004 mayoral election.

To assess whether adopting the reform is associated with hospital characteristics that trend differently (e.g., hospitals that are performing better over time are more likely to adopt the new recruitment system), column (3) presents the OLS coefficients of a regression of the first difference of each characteristic on a dummy that takes value 1 for ever-treated hospitals and 0 otherwise. We do not observe that treated units are on significantly different trends than never-treated hospitals, in terms of patient composition, hospital outcomes, or political determinants. We consider these results as supporting evidence for our research design.

Main Results

We begin by estimating the following staggered difference-in-differences model (DiD):

$$y_{ht} = \alpha_h + \gamma_t + \beta \times \mathit{Reform}_{ht} + X'_{ht}\Delta + \epsilon_{ht}, \tag{2.1}$$

where y_{ht} is an outcome variable at hospital h and quarter t level, and Reform_{ht} is a dummy variable that takes value 1 if a hospital adopts the new selection process and 0 otherwise. Recall that once a hospital selects a CEO via the new recruitment system, it has to select all future managers using the same recruitment system. Thus, adoption of the recruitment reform is an absorbing treatment and the dummy variable takes the value 1 for all periods after the first manager is hired under the new regime. The control group consists of yet-to-be treated and never-treated hospitals. α_h represent hospital fixed effects that control for unobservables specific to the hospital and γ_t are time fixed effects to account for unobservable shocks specific to a quarter.

To account for differences on patient characteristics, we follow the literature and include X_{ht}, a comprehensive set of hospital-by-quarter level variables that pick up differences in case mix characteristics (Propper and Van Reenen, 2010; Gaynor et al., 2013). Specifically, the vector X_{ht} includes the share of female inpatients, the share of inpatients within each of eight age bands, the share of inpatients within each of the 31 categories of the enhanced Elixhauser comorbidity index (Elixhauser, Steiner, Harris, and Coffey, 1998; Quan, Sundararajan, Halfon, Fong, Burnand, Luthi, Saunders, Beck, Feasby, and Ghali, 2005), and interaction of the various shares. To control for the socioeconomic status of patients, X_{ht}

also includes the share of inpatients with each type of health insurance. We cluster standard errors at hospital level, which is the treatment-level unit. The coefficient of interest is β, which summarizes the impact of the reform on hospital quality.

For estimation, we consider the universe of public hospitals and weight each regression by the number of inpatients as of 2005.[19] In Table 2.4 we report the $\hat{\beta}$ obtained from estimating Equation 2.1 using different death-related measures of hospital performance. Columns (1)-(3) consider the logged in-hospital death rate. Column (1) shows that reform adoption led to a 13% decrease in in-hospital death rates, and columns (2)-(3) confirm that the result is robust to adding the set of case mix characteristics discussed above, either separately or interacted. Column (4) considers the log of the death rate 28 days after admission, and includes both in- and out-of-hospital deaths. Reassuringly, the point estimate shows that effects are not driven by higher out-of-hospital deaths. Column (5) focuses on emergency admissions, which should be less prone to non-random sorting, and finds a similar impact of the reform in this sample. Finally, in column (6) and (7) we use a Poisson model to estimate the effect of the policy on the number of deaths. Column (6) focuses on all deaths and shows that the reform decreased deaths by around 5.7% (i.e., $\exp(\hat{\beta})-1$, where $\exp(\hat{\beta})$ is the incidence rate ratio of deaths). Column (7) shows that the death rate among emergency cases with acute myocardial infarctions (AMI, commonly known as "heart attacks") decreased by around 14.6%, although this coefficient is more imprecisely estimated.[20]

Validity of Results and Alternative Explanations

In this subsection, we discuss the validity of the above results. We first present event study evidence that provides visual support for the assumption of parallel trends. Next, we discuss whether patient selection could be driving our results. Finally, we examine whether CEO transitions have, per se, a mechanical effect on hospital quality.

Testing for Parallel Trends: Event Study Evidence

We start by assessing the existence of pre-trends. The concern here is that hospitals that adopt the selection reform might be on different trends to those that have not adopted it, which could bias our results. To partially assess the validity of the underlying parallel trends assumption, we estimate the following event study:

$$y_{ht} = \alpha_h + \gamma_t + \sum_{k=-6}^{12} \beta_k D_{ht}^k + X_{ht}'\Delta + \epsilon_{ht}, \tag{2.2}$$

[19] For hospitals that had a CEO turnover, we include a window of 6 quarters before and 12 quarters after reform adoption to facilitate study of the timing of the effect.

[20] Note that the number of observations drops from 8,104 to 1,956. Following the literature, we define AMI deaths as deaths that occurred 28 days after admission of patients (through the emergency room) with an ICD 10 diagnosis of I21 (Acute myocardial infarction) or I22 (Subsequent myocardial infarction). For estimation, we weight this regression by the number of inpatients with emergency room AMI admission as of 2005.

where D^k_{ht} is a dummy variable that indicates the reform was adopted k periods earlier (or will be adopted k periods ahead for negative values of k). Reform adoption is an absorbing treatment. The β_k coefficients can be interpreted as the effect of the reform on hospital quality for each k quarter, relative to the quarter before adoption. We normalize the coefficients such that $\beta_{k=-1} = 0$, and we consider a window of 6 quarters before and 12 quarters after adoption.

Figure 2.2 displays the point estimates of our β_k and their confidence intervals for different measures of hospital-level death rates. When inspecting the dynamic effects of reform adoption, we observe that—across all panels—the pre-period estimates tend to be small, around zero, and not significant, which indicates that treated and control units were not on different trends prior to reform adoption. Furthermore, after the reform, the estimates turn negative and significant and increase gradually. In this case, it does not seem that the change in management is driven by a previous worsening (improvement) in managerial performance, which would overestimate (underestimate) the true impact (if any) of the treatment. In Appendix Figure B.6.5, we conduct robustness checks and plot the impact of the reform on the count of in-hospital deaths using a dynamic Poisson model (Panel A), and on the log of predicted death rates and on the ratio of actual over predicted death rates using a two-way fixed effects model (Panel B). Finally, in Appendix Figure B.6.6 we present estimation results using the models suggested by De Chaisemartin and d'Haultfoeuille (2020) and Borusyak, Jaravel, and Spiess (2022a), which are robust even if the treatment effects are heterogeneous over time or across groups. Results are robust and follow the same dynamic trajectory regardless of estimation strategy.

Risk-Adjusted Results

To ease selection concerns, we also estimate the impact of the reform on the actual over the predicted death rate and only on the predicted death rate as a placebo exercise. We follow the prediction procedure used by the CMS (Ash, Fienberg, Louis, Normand, Stukel, and Utts, 2012), which predicts the likelihood of death at the individual level on a detailed set of patient characteristics. We first fit a logit model for the outcome of death using the set of case mix controls and more than 5.5 million patient-level observations from 2001 to 2004. Then, we use the model's predicted death probability for each patient (based on patient case mix) to obtain "predicted" death rates at the hospital level. Finally, to ease interpretation and following the UK's NHS (Health and Centre, 2015), we construct a "risk-adjusted mortality rate" that divides the actual hospital-level death rate by the predicted death rate, such that an increase (decrease) from one means more (fewer) deaths than predicted deaths.[21]

Table B.6.2 shows the robustness of our result to using alternative risk-adjusted measures. Columns (1)-(3) show estimates from Equation 2.1 obtained for different definitions of the risk-adjusted death rate. In Column (1), the risk-adjusted death rate is based on patients' demographics (gender and age). Column (2) also considers patients's health insurance, a

[21] We provide further details of the the risk-adjustment procedure in Appendix B.3.

proxy for socioeconomic status. Finally, column (3) corresponds to our preferred measure that also includes patients' diagnoses based on the enhanced Elixhauser comorbidity index. We find that the policy had no effect on the predicted death rate, but a significant impact on the ratio. After the new CEO selection process was adopted, the ratio of actual over predicted death rates decreased by 8% from a base of 0.79 in our estimation sample. Results are stable regardless of the incorporation of more covariates in the logit model. This is reassuring because, according to recent research that leverages quasi-random variation on death rates, risk-adjusted mortality measures are reliable and valid indicators of hospital quality in the U.S., where the institutional setting is prone to patient selection (Doyle et al., 2019).

Testing for Patient Selection

The risk-adjustment procedure is fundamentally based on patient diagnoses, which raises three potential concerns. First, new managers may have incentives to influence the diagnoses for billing or revenue purposes (Silverman and Skinner, 2004). Second, new managers may reject sicker patients based on the severity of their illness. Finally, there could be substantial variation in diagnostic practices across doctors and regions unrelated to patients' characteristics.[22] If, for example, managers bring in new doctors who, in turn, have systematic differences in diagnostic practices, our results could be explained by a mechanical effect of doctor composition.

Careful consideration of our setting's characteristics suggests that the first two concerns are unlikely to drive our results. On the one hand, the diagnoses in our data come from a nationwide mandatory program that aims to characterize the morbidity profile of patients for policy purposes and are recorded directly by the lead physician (Decreto 1671 Exento, 2010). Therefore, there is no clear way the hospital CEO could manipulate diagnoses. On the other hand, the law forbids CEOs from selecting patients based on their condition and must adhere to referral and counter-referral guidelines. Furthermore, we can empirically assess these three concerns by examining whether hospital-level risk-score changes upon a new manager appointment. For this purpose, we use the patient-level data to fit a logit model of (pre-reform) mortality on patients' demographics and diagnoses (for details, see Appendix B.3). Then, we predict the probability of death for each patient, and use these predictions (i.e., patient level risk scores) to construct hospital-level predicted death rates and number of deaths. We estimate Equation 2.1, but now replace the dependent variable with these predictions. Figure B.6.7 plots event study evidence on the null effects of the reform on mortality predicted based on patients' risk-scores.

Although our results are robust to adding case mix controls and using risk-adjusted mortality measures, there could be selection on unobservable patient characteristics that is not picked up by diagnosis data or the list of available patient characteristics. For instance, perhaps managers are able to reject sicker patients in a way that does not change patient

[22] See, for example, Song, Skinner, Bynum, Sutherland, Wennberg, and Fisher (2010) and Finkelstein, Gentzkow, Hull, and Williams (2017), who document and discuss this phenomenon in the United States.

composition (supply-side selection on unobservables), or healthier patients are more likely to go to a given public hospital if they observe that public hospital performance is improving (demand-side selection on unobservables).

To indirectly test whether supply-side selection on unobservables lead our estimates to be upward biased, we consider the impact of the reform on mortality rates in nearby hospitals and deaths at home. To the extent that rejected patients die, they would still show up in the mortality statistics of the hospital's geographic area. For this exercise, we estimate Equation 2.1 again but now use as dependent variables the at-home death rate (in the municipality where each hospital is located) and the in-hospital death rate in nearby hospitals. Panel A in Figure 2.3 shows our results, with baseline estimates as a reference. We find that adopting the reform in a given hospital has no significant impact on at-home death rates in the hospital's municipality or on the death rates of nearby hospitals.

Finally, to examine whether sorting on unobservables is biasing our results, we exploit two features of our setting. First, we leverage the fact that a set of individuals are locked-in in the public health network, i.e., lower-income individuals receive public insurance for free but cannot use it in private providers. Second, we can use our data to empirically identify the set of patients that comply with the referral guidelines described in Appendix B.1. For this analysis, we estimate Equation 2.2 using smaller samples comprised exclusively of locked-in patients and referrals compliers patients. The results from this approach—which should mute demand-side sorting if any—are presented in Figure 2.4. Reassuringly, we in both restricted samples we find a similar impact of the reform on hospital performance.

Manager Transition Mechanical Effect

We next examine the extent to which there is a mechanical effect on death rates due to the CEO transition itself. For instance, an alternative explanation for our results could be that the decline in the death rate reflects the effect of the arrival of a new manager, by means of a Hawthorne effect (Acemoglu et al., 2022).

Exploring this mechanism requires slightly modifying our empirical strategy, since all hospitals have several transitions in the examination period. To deal with multiple events and the lack of clean controls, we perform a stacked event study (Cengiz, Dube, Lindner, and Zipperer, 2019a; Baker, Larcker, and Wang, 2022; Atal, Cuesta, González, and Otero, 2022a). We define an event as a CEO transition in a never-treated or yet-to-be-treated hospital in any quarter between 2001 and 2019. For each transition event, we define a time window around it and a control group of hospitals with no transitions in the time window.[23] Next, we define a set of valid events as those that are balanced in the time window and do not overlap with another transition in the pre-period within the time window. Finally, we

[23]Note that there is a trade-off between the length of the window and the number of events and controls. We use 4 quarters prior to the transition and 12 quarters after the transition, although the results are robust to other time windows.

append the data for all valid events and estimate the following equation:

$$y_{hte} = \alpha_{he} + \gamma_{te} + \sum_{k=-4}^{12} \beta_k D_{hte}^k + \epsilon_{hte}, \tag{2.3}$$

where e is a valid transition event. Equation 2.3 is analogous to Equation 2.2, but the observation is at hospital-by-time-by-event level and replaces the hospital and time fixed effects with hospital-by-event and time-by-event fixed effects. We cluster standard errors at hospital level.

Figure 2.5 presents the effect of a CEO transition on death rates. The effect is a precisely estimated zero and confirms that a CEO transition before the reform has no significant effect on hospital quality. This evidence suggests that the impacts of the recruitment reform reported so far are not explained by a mechanical effect driven by the CEO transition itself.

How is the Reform Improving Hospital Performance?

Now we ask what are the underlying mechanism that explain the drop in mortality that we observe after the adoption of the reform. We focus on the efficient use of hospital resources and personnel practices. Most data used for this analysis is only available at yearly frequency. Thus, we perform the analysis at this level of aggregation. For completeness, in Appendix Figure B.6.8 we present the effect of the reform on mortality rates using this level of aggregation.

Operating rooms (ORs) are one of the most critical hospital resources and typically account for more than 40% of total expenses (Association et al., 2003; Denton, Viapiano, and Vogl, 2007; Guerriero and Guido, 2011). Inefficient use of ORs is extremely costly for patients and can impact overall hospital performance.[24] The efficient use of ORs is a highly complex operational and management problem, and management practices are a crucial lever for improving OR efficiency (see, e.g., He et al., 2012).[25] Another critical element for hospital performance is high-skilled personnel. Recent literature in personnel economics shows that better-managed firms recruit and retain workers with higher human capital (Bender et al., 2018). These reasons lead us to examine OR efficiency and personnel

[24] Late starts, or longer-than-expected surgeries trigger delays or rescheduling for patients next in line. In turn, to deal with surgeries that finish after their rostered times, the medical staff has to work overtime, which implies direct costs to the hospital and can lead to higher levels of burnout, medical errors, and patient dissatisfaction (Rogers, Hwang, Scott, Aiken, and Dinges, 2004; Denton et al., 2007; Stimpfel, 2012). The other main effect of the inefficient use of ORs is that hospitals can treat fewer patients, and hence patients face longer waiting times (Durán, Rey, and Wolff, 2017).

[25] For instance, planning and scheduling must consider OR availability and match the workload to medical staffing, the material resources required, and the availability of post-surgical recovery beds (Wang, Demeulemeester, Vansteenkiste, and Rademakers, 2021). Furthermore, OR planning and scheduling must incorporate the uncertainty entailed in surgery duration and emergent admissions that require a surgical procedure (Latorre-Núñez, Lüer-Villagra, Marianov, Obreque, Ramis, and Neriz, 2016).

practices as mechanisms through which new managers improved hospital performance after the reform.

We find that the recruitment reform had a significant and economically meaningful effect on OR efficiency. We run the same specification as in Equation 2.2 on the logged ratio of OR utilization (i.e., the number of hours ORs are used) to OR capacity (i.e., the aggregate available number of OR hours). Panel A in Figure 2.6 shows the effect of the reform on the ratio between OR utilization and capacity. We find that the reform did change the number of hours ORs are effectively used: 3 years after the reform adoption, the number of OR hours used increased by 25%. Although this number might seem big at face value, it is not large enough to close the gap between the average efficiency in high-complexity hospitals and the average in the UK's NHS.[26]

Panel B in Figure 2.6 examines the other side of the coin of higher OR usage, the number of surgeries performed. We find that the number of surgical procedures increased by a magnitude similar to utilization of ORs (yellow diamonds). To deepen this analysis, we follow an agnostic procedure to classify diagnoses as amenable to death prevention through surgery. Specifically, we use our patient-level records to estimate several logit regressions of (pre-reform) mortality on surgery indicators while controlling for patients' demographics. We estimate one regression for each of the 31 categories of the enhanced Elixhauser comorbidity index Elixhauser et al. (1998); Quan et al. (2005) and then based on the estimated coefficient for the surgery indicator, we classify the diagnoses into amenable to death prevention through surgery if $z \leq -2.576$. As shown by Figure 2.6, Panel B, we find that most of the increase in surgeries was led by a surge in surgeries related to diagnoses amenable to death prevention through surgery (e.g., Cardiac Arrhythmia, Renal Failure, Metastatic Cancer).

We examine the reform's effects on personnel turnover and wages. For this, we use administrative data on hospital personnel coming from the Human Resources Information System used for HR purposes by the Ministry of Health. We run the same specification as in Equation 2.2 on the likelihood that a worker will leave the next period (either job to job or job to unemployment transitions) and on their logged hourly wages. Panels C and D of Figure 2.6 show our results. We find that the reform reduced the turnover of doctors, but it did not change wages, which is expected given that in the public sector, wages are rules-based. From anecdotal evidence based on conversations with managers and doctors in the public sector, we posit that the reduced turnover rate might be explained by unobservable benefits and amenities that the manager can negotiate with doctors.

CEO Selection Reform in the Context of Other Policies

We conclude this section by benchmarking our results to the effects of other policies studied in the literature. One of the advantages of our data is that we can check the impact of the policy on different samples of patients, which allows us to match some of the characteristics

[26] Out of 9 hours of daily capacity, the average in a sample of high-complexity hospitals in Chile is 4.8 hours and in the NHS is 6.4 hours (CNEP, 2020).

in the sample of patients studied elsewhere. For each comparison, we present the average death rate in the sample in the literature and in ours after we match it according to patients' characteristics. Note, however, that although we can match the sample of patients in, say, age-bracket and type of admission, patient composition will still differ across settings. Comparisons should thus serve as a benchmark and not as a horserace competition between policies. Results are summarized in Table 2.5.

We start by comparing the effect of the CEO selection reform with the impact of increasing health spending. Doyle et al. (2015) examine the effect of receiving higher payments from Medicare. They find that a 10% increase in Medicare reimbursement per capita decreases death rates by 6%. Their sample of patients includes emergency admissions arriving by ambulance, over 65 years old, and with non-deferrable medical conditions. Since we do not have records on whether a patient arrives by ambulance, we only compute the effect of our policy on the sample of patients over 65 admitted via the ER. We find a similar effect over a very similar average death rate in the sample.

As a second comparison, we consider recent evidence on the impact of public vs. private provision of healthcare. Chan et al. (2022) study the case of VA hospitals in the US and find that public provision reduces 1-year mortality by 7.7% in veterans over 65 years admitted from an ambulance. We find a similar effect in the sample of emergency admissions over 65 years. As noted above, we cannot observe whether a patient is arriving by ambulance. Nonetheless, we find a very similar effect size over a very similar average death rate.

Finally, we focus on policies related to the impact of increasing competition in the health sector. Bloom et al. (2015b) examine the effect of adding competition between health providers in the UK. They find that adding one extra hospital in the neighborhood decreases the in-hospital 28-day death rate by 10% following emergency admissions for AMI. The policy we study in this paper finds a similar effect, although over a higher average death rate in the same sample group. Previous work by Gaynor et al. (2013) also reports that increasing competition by 10%, as measured by a decrease in the Herfindahl-Hirschman Index (HHI), reduces the 28-day in-hospital death rate by 1% and the overall death rate by 20%. In this regard, improving CEO selection has a comparable effect of 15% and 20% reduction in deaths rates, but over a much larger sample mean.

2.4 Reform Financial Incentives Effects

The reform included a change in the recruitment system and financial incentives, in the form of pay for performance and higher base wages. Low-powered incentives and low wages in the state are often pointed to as one source of the inefficient performance of high-end public employees. For instance, recent empirical research shows that financial incentives can increase the performance of employees in the public sector (Khan, Khwaja, and Olken, 2015; Biasi, 2021; Deserranno, Caria, Kastrau, and León-Ciliotta, 2022).

Perhaps post-reform managers improved hospital performance simply because they exerted more effort due to the newly introduced financial incentives. In this section, we examine the financial incentives effects of the reform and find that our results are not explained by either performance pay or higher wages.

Results Are Not Driven by Performance Pay

According to performance-related pay models, performance pay incentives attract higher-ability workers and also induce them to exert greater effort (Lazear, 2000). In our setting, the head of the Health Service (i.e., the principal) writes a performance contract in agreement with the hospital CEO (i.e., the agent) for a 3-year period. At the end of each year, the CEO gets a final score based on the parameters in the contract. The yearly wage is impacted by the performance agreement according to the following schedule:

$$\text{Yearly Wage}_t = \begin{cases} 100\% & \text{if} \quad \text{performance}_{t-1} \geq 95\% \\ 98.5\% & \text{if} \quad 65\% \leq \text{performance}_{t-1} < 95\% \\ 93\% & \text{if} \quad \text{performance}_{t-1} < 65\%. \end{cases} \tag{2.4}$$

Two things are worth noting about the schedule in Equation 2.4. First, the wage in the first year is not affected by the schedule because it is based on the previous year's performance, and the performance pay penalty only affects years 2 and 3 of the agreement. Second, the reform introduces only a small penalty and no possibility of a wage increase. The maximum penalty is a 7% discount of the yearly wage.

We accessed all available performance contracts for the first manager appointed after the reform adoption and their yearly performance scores.[27] Figure 2.10 presents the distribution of the 3-year average performance score. Note that 60% of the distribution is above the 95% threshold and avoids any wage penalization. The rest are between 95% and 65%, which is the lowest threshold to avoid a 7% wage penalty. No manager receives a score below the 65% performance threshold. This evidence suggests that the performance agreements were not binding, and most managers easily met performance goals. In Appendix B.4, we empirically analyze whether CEOs' performance scores implied better managerial performance at the hospital. We find that managers with high and low performance scores were equally effective in improving hospital performance.

We note that the performance agreements included in the recruitment reform were poorly designed across the board, and their lack of effectiveness is not specific to public hospitals. For example, in all government positions that used the recruitment system, less than 5% scored less than 80% on their performance scores in 2013 (CPPUC, 2013), and more than 90% achieved a 100% performance score in 2016 (CADP, 2017). This tool's failure to be a useful management control has been addressed in several policy reports calling for its amendment (see, e.g., CPPUC, 2013; Barros, Weber, and Díaz, 2018). We conclude that in our setting, performance pay is not likely to be a relevant driver of managerial productivity.

[27] Unfortunately, some of the oldest contracts and performance scores are lost, and the Civil Service has no available records. Out of 87 processes, we have performance data for 57 and access to 77 contracts.

Results Are Not Driven by Higher Wages

An alternative mechanism is that the results are driven by efficiency wages. According to this hypothesis, wages above their outside option create an incentive for managers to exert extra effort and can elicit productivity growth (Katz, 1986). If the reform bonus creates labor rents, then this mechanism might be at play.

We start by analyzing the reform bonus. The bonus consists of an increase in the base salary, which is defined for each position by the Ministry of Finance. We document the size of the reform bonus relative to the position's pre-reform pay in two ways. First, in Appendix Figure B.6.10 we present a box plot of the share of the quarterly remunerations that is explained by the reform wage bonus. The reform bonus explains, on average, 43% of the quarterly wage, and the middle 50% of the distribution is between 37%and 46%. In Appendix B.5, we present event study evidence on the reform's effect on CEO wages when it is adopted by a hospital. On average, we find an effect of the same order of magnitude, albeit somewhat smaller. However, it is important to note that we do not observe the change in the CEO's remuneration but rather in the position's remuneration. Hence the effect is a composition of mechanical changes in pay due to changes in the manager's identity and the pay increase.

To examine the potential effects of efficiency wages in this setting, we exploit a 2016 amendment to the law that created the recruitment reform (Ley 20,955). Among other things, the amendment changed the pay scheme in the following way. Before the amendment, all CEOs were paid according to the public employees' pay grade, regardless of their profession. After the modification, CEOs appointed after November 2016 can choose to be paid according to the medical pay laws instead of the public employees' pay grade *only if* they are doctors.[28] The medical pay law is much more generous than the public employees' pay law. Therefore the amendment implied an increase in remuneration for doctor CEOs but not for CEOs with other backgrounds.

If the efficiency wage hypothesis is at play in this setting, we should expect that a wage increase is followed by an improvement in performance in hospitals in which new managers are doctors *and* receive a pay boost. To study this question, we perform a stacked event study, in which an event is a transition after November 2016 that uses the new selection system and the incoming CEO is a doctor. For each event, we define a time window around the transition and determine an event-specific control group that includes hospitals with no transition and units with transitions to professionals other than doctors. We select valid events that are balanced in the time window and that do not overlap with other transitions one period before the event.[29] We then append the data for all valid events and estimate an event study following Equation 2.3.

Panels A and B in Figure 2.11 present the impact of the 2016 amendment on doctor CEO wages and hospital performance, respectively. As expected, the change in the regulation

[28] More precisely, doctors can choose to be paid according to Law 19,664 instead of Law 18,834.

[29] As noted before, there is a trade-off between the length of the window and the number of valid events. In total, there are 33 events and 24 valid events.

increased wages for incoming doctor CEOs. The effect is around a 10% quarterly wage increase. However, we do not observe any effect on death rates. In other words, the wage increase was not followed by an improvement in CEO performance. This finding suggests that the efficiency wage hypothesis is unlikely to play a substantial role in this setting. Therefore, we rule out this hypothesis as a significant driver of our main results.

All in all, the evidence suggests that financial incentives do not explain the performance improvement we observe after adoption of the selection reform. In other words, had the reform only included financial incentives and not changed the recruitment system, we do not expect to observe an impact on hospital performance. Therefore, the key component of the reform was the introduction of the competitive recruitment system, which changed the identity and characteristics of hospital CEOs. We examine this mechanism in the next section.

It is important to note, however, that although we have ruled out the possibility that financial incentives play a role in managerial performance, this result is conditional on the selected CEO. The extra pay likely plays a role in the decision to apply. For instance, Dal Bó et al. (2013) show that higher pay for public sector positions attract more competent applicants. Unfortunately, we do not have a design to test this hypothesis because we do not observe the pool of applicants *before* adoption of the recruitment reform in each hospital. It is an open question to what extent higher wages widen the pool of high-quality applicants in our setting and, through this mechanism, higher wages impact performance. For instance, it could be the case that appointed CEOs with management studies would have been less likely to apply in the absence of the wage hike.

2.5 The Recruitment Effects of the Reform

Impact of the Reform on CEO Characteristics

The evidence so far suggests that the impact of the reform on hospital performance is not driven by the financial incentives in the reform. In this subsection, we examine the recruitment effects of the policy and evaluate how the new recruitment process changed the characteristics of new CEOs. To this end, we use the same research design as before but replace the independent variable with manager-specific characteristics. Concretely, we estimate Equation 2.1 on $X_{M(h,t)}$, where X are individual-specific traits such as educational background, skills, and demographics, and $M(h,t)$ is a function that indicates the identity of the CEO of hospital h at time t.

We start by computing the impact of the policy on educational background. We focus on management, which is one of the targeted backgrounds of the policy. We measure management studies using two complementary variables. First, we construct a variable that takes the value 1 if the individual has an undergraduate degree with management coursework and 0 otherwise. We consider the following majors to include management courses: public administration, business and economics, accounting, and engineering. The second variable relates to postgraduate education in management. This variable takes the value 1 if in a given quarter an individual has postgraduate management studies and 0 otherwise. Postgraduate management studies include master's degrees and diplomas related to management and administration. For example, the former include master's degrees in public health administration, public administration, and business administration, among others. Diplomas are shorter executive education courses, akin to professional certificates in the US.

Figure 2.7 presents the results. Panel A shows that the reform increased the share of CEOs with undergraduate management degrees by 21 percentage points, from a baseline of only 2%.[30] The increase in the number of CEOs with this background came at the expense of displacing almost one-to-one doctor CEOs, who before the policy adoption accounted for 93% of CEO positions, and a slight negative effect on health professionals other than doctors. The reform did not have an effect in hiring professionals with a background in other disciplines. Importantly, Panel B in Figure 2.7 shows that the displacement of doctor CEOs masks heterogeneous effects. In fact, the policy increased the number of doctor CEOs with postgraduate management studies by 14 percentage points—from a baseline of 18%—while substantially decreasing the number of doctor CEOs without management studies by 31 percentage points—from a baseline of 75%.

To further examine this heterogeneity, in Figure 2.8 we plot the the dynamic effects of the policy in a 3-year window after adoption. An interesting finding is that the displacement effect on doctor CEOs significantly wanes over time. Panel A focuses on the likelihood that the CEO has a management undergraduate degree or a medical degree. The reform increased CEOs with a management undergraduate degree by around 25 percentage points the quarter immediately after adoption, but the effect decreased over time to slightly less than 15 percentage points.[31] In the case of doctors, we observe the opposite effect. After adoption, there is an initial displacement of around 20 percentage points. But by the end of the 3-year window, doctors were able to revert the loss in their likelihood of securing a CEO position to only 10 percentage points. Panel B in Figure 2.8 decomposes the total effect on doctor CEOs into the change coming from doctor CEOs with management training and doctor CEOs with no management training. While on impact the reform had a negligible effect on the likelihood that doctor CEOs have postgraduate studies related to management, by the end of the 3-year window the effect increased up to 20 percentage points—more than duplicating the pre-reform average. The flip side is that the policy decreased permanently—and on impact— the share of doctor CEOs with no management studies by around 30 percentage points.

In Figure 2.9 we study the impact of the reform on the likelihood that the CEO had completed *any* management studies before her appointment. The average across ever adopters 1.5 years before the reform was 21%.[32] The reform increased the likelihood that the CEOs holds

[30]Since the timing of adoption varies across hospitals, we compute the baseline in the period before each hospital adopted the reform.

[31]The effect changes over time because CEOs' tenure is, on average, shorter than 3 years, and therefore the effect is picking up the characteristics of more than one post-policy manager.

[32]As a point of comparison, in NHS hospitals, (Janke et al., 2020) report that 26% of CEOs hold postgraduate

a management undergraduate degree or management postgraduate studies by 37 percentage points. Importantly, the effect is stable over time and is explained both by professionals with management undergraduate degrees and doctors with management training taking over CEO positions after the reform.

The increase in the share of doctor CEOs with management studies is a combination of two phenomena. First, the reform likely increased the chances of being appointed CEO for the pool of doctors who would have had management studies in the absence of the reform. But the policy also incentivized doctors who wanted to be appointed CEOs to pursue formal management studies in order to improve their competence and the likelihood of passing the recruitment process. The reason is that securing a CEO position is more likely if the candidate has management studies. As the reform was gradually adopted across the public health sector, management studies were also more demanded by doctors. This second explanation is consistent with the fact that before 2003—the year when the reform was enacted—there was no supply of master's degrees in health management. Indeed, as shown in Appendix Figure B.6.9, the timing of opening of health management postgraduate programs coincides with the timing of the reform we study. The figure also shows that management postgraduate programs focused on areas other than health were available for a long time before and gradually increased over time. Qualitative anecdotal evidence further supports the claim that these new programs are geared toward doctors seeking careers in health administration.[33]

In Table 2.6, we further investigate the impact of the policy on other CEO characteristics. We use the estimated CEO fixed effects from Appendix B.2 as a measure of managerial skills, and the performance on the standardized national university entrance exam as measure of cognitive ability. For demographics, we consider age and gender. Column (1) shows that the reform led to the appointment of CEOs with higher managerial talent—as proxied by their CEO fixed effect. Column (2) shows that the new managers performed slightly worse on college entrance exams, although this difference is not significant. Considering that the new managers are displacing doctor CEOs and that to secure a position in medical school individuals need to achieve top scores on college entrance exams, this result implies that on average, new managers are also top performers in college entrance exams. In columns (3) and (4), we focus on the set of CEOs who took the older version of the college entrance exam in Chile, in which applicants had to choose which *specific* exam to take. We find that new managers are more likely to take the math-specific exam and less likely to take the science exam. This finding is consistent with the fact that the reform increased the share of CEOs with management-related undergraduate degrees.

Columns (5) and (6) focus in demographics. We find that the new managers are on average almost 2 years younger than the CEO would have been in the absence of the policy.

managerial education. Bloom et al. (2020) provide an additional antecedent and document that in a sample of hospitals in nine developed and developing economies, on average only a quarter of managers (including non-CEO managers) report having received management training.

[33] See, for example, this news report as a case study: https://www.americaeconomia.com/articulos/notas/mba-en-salud-para-que-medicos-chilenos-entren-al-mundo-del-management.

One interesting finding is that the reform did not have any impact on female participation in CEO positions. The average pre-policy share of female CEOs is 22%, which is in line with the widely documented under representation of female CEOs in the private sector; this phenomenon is known as the glass ceiling (Bertrand, 2018). We find that in this setting, the reform had no discernible effect on the likelihood of women making it to the top. This is also consistent with recent research showing that the application of unwritten impartial hiring processes in the public sector does not have an effect on gender hiring disparities (Mocanu, 2022).

Could the Attenuation of Skills Mismatch Drive the Results?

We now ask which factors can explain the effectiveness of the new managers. In particular, we examine the extent to which new managers are higher performers due to a better match between their skills and the skills demanded by the job. In the public sector, several factors may create skill mismatches that may hinder performance.[34] In the case of public hospitals, as discussed below, the social norm before the reform was that CEO positions were reserved for doctors. The reform mitigated the skills mismatch by displacing doctor CEOs for professionals with management degrees and also incentivized doctors who wanted to pursue careers as hospital CEOs to invest in management education. We examine whether correcting the skills mismatch in this setting enhances the organization's performance.

Concretely, we refer to skill mismatch as the extent to which individuals are employed in an occupation unrelated to their main field of study. This phenomenon is known as horizontal mismatch, as opposed to vertical mismatch, in which individuals have a higher or lower educational attainment than needed for their job. While a nascent literature studies horizontal mismatch in the private sector, to the best of our knowledge there is limited or no research in the public sector (Nordin, Persson, and Rooth, 2010; Besley et al., 2022).

To examine whether CEOs with management studies perform better than those without, we interact the reform dummy in Equation 2.1 with a dummy that takes value 1 if the CEO has management studies and 0 otherwise. The working assumption is that CEOs with management studies are well matched, while the rest represent mismatches.[35]

Columns (1)-(3) in Table 2.7 present the results. In column (1), we find that when the appointed CEO has management undergraduate studies, the policy has a larger effect than when she does not. The point estimate for matches is larger than the effect for mismatches, although the difference is not significant at standard confidence levels. In column (2), we use a less demanding definition of mismatch and compute the differential effects of the policy in

[34] For instance, a combination of low exit rates among public employees and technological change (Besley et al., 2022).

[35] One limitation of this analysis is that due to a lack of data, we neglect heterogeneities in management experience for individuals without formal studies in management. Implicitly, by abstracting management experience from the analysis, we assume that any management skill acquired in the job is firm-specific, while skills acquired from formal education are general management skills and are transferable across units (Frydman, 2019).

the cases in which the manager has *any* management studies, including undergraduate and postgraduate studies. We find that now the difference is starker and statistically significant at a 99% confidence level. In column (3), we focus only on the sample of CEOs who are doctors. The interaction captures the differential effect of the policy between doctor CEOs who have management studies and those who do not. We find that the policy had a significant effect in cases in which appointed doctor CEOs had management studies and negligible effects otherwise.

Another way to examine treatment effect heterogeneity based on management studies is to compare transitions from CEOs *without* any management studies to CEOs with management studies (as of the time of the transition). As before, management studies refer to undergraduate and postgraduate studies related to management. Concretely, we estimate Equation 2.3 in an stacked event study framework. We define an event as a CEO transition, and select a set of *valid* events that are balanced in the time window and do not overlap with other transitions for at least four periods before the event. For each event, we define a time window around a transition event and a control group of hospitals with no transitions in the time window.[36] To avoid a mechanical correlation with the results presented in columns (2) and (3) in Table 2.7, we exclude all CEO transitions that occurred the first time the selection reform was implemented in a given hospital. Finally, we append the data for all valid events and estimate an event study.[37]

Columnns (4)-(5) in Table 2.7 present the results. Column (4) presents the 3-year effect of a CEO transition without any management studies to a CEO with management studies. Column (5) is a placebo exercise that estimates the effect of transitions between CEOs *without* management studies. Consistent with the findings in the interacted DiD, we find that when the hospital transitions from a CEO without any management studies to a CEO who has management studies, the event is followed by a significant decrease in death rates. We find no effect on hospital death rates when we examine transitions between CEOs without management studies. Both effects are also consistent with the evidence presented in columns (2) and (3) for the effect of the reform when the appointed manager had and did not have management studies.[38] This evidence suggests that hospitals transitioning to CEOs with management studies drive, for the most part, the effect of the selection reform on hospital performance.[39].

Interestingly, the finding that management studies improve CEO performance might be

[36] As noted before, recall that there is a trade-off between the length of the time window and the number of valid transitions and control units.

[37] Appendix Table B.6.4 presents the number of events by type of transition.

[38] Appendix Figure B.6.14 provides visual event study evidence of the effect of transitions from CEOs without management studies to CEOs with management studies. Importantly, we find no pre-trends and the same trajectory as the effect of the reform displayed in Figure 2.2. The lack of pre-trends suggests that the hospital's performance does not drive the change in the education of the CEO.

[39] In Appendix Table B.6.6 we regress the CEO managerial talent—measured by the CEO fixed effect estimated in Appendix B.2—on CEO observable characteristics. These characteristics include gender, age, age squared, and a set of indicators for their educational background. Consistent with the findings above, management studies are correlated with better managerial performance (i.e., lower CEO fixed effects)

at odds with the results of Acemoglu et al. (2022), who show that managers with a business degree do not improve their firm performance and reduce their employees' wages by means of rent-sharing practices.[40] One key difference is that in our setting, business managers perform in the public sector, in which they have fewer incentives to reduce their employees' wages and fewer ways to do so, given the rigidity of public sector wages. Furthermore, business CEOs who self-select into the public sector might have higher levels of prosocial motivation than those in the private sector (Finan, Olken, and Pande, 2017).

What explains skills mismatch in public hospital CEO positions?

Given the significant impact on performance delivered by CEOs with management training, why are not all public hospitals run by CEOs with this background? A primary reason is that before the implementation of the policy, there was a strong social norm in the public health sector that hospital CEO positions were reserved for doctors. Although there is no statutory rule prohibiting non-medical professionals from being selected as CEOs, in 2004—the year before the first hospital implemented the selection reform—doctors made up 98% of public hospital CEOs. The policy had a big effect on changing this empirical fact: By 2019, the fraction of CEOs with medical degrees in treated public hospitals was 53%.[41]

Anecdotal evidence allows us to understand why this norm emerged and was sustained over time. According to the responses in a small survey conducted by Civil Service on public hospital CEOs, doctors tend to believe that individuals with no medical training should be restrained from CEO positions. For instance, the view of one doctor CEO was that *"the ideal place for the engineer is as an advisor to a doctor CEO. The engineering vision is super positive and necessary for organizing finances, indicators, goals, etc., but they have a very large information asymmetry with the medical team. A doctor can tell the non-medical CEO 'you don't understand this, you can't comment' and that's it."* (Servicio Civil, 2014).[42]

This same belief may have discouraged doctors from investing in management training. If doctors thought management training would not improve their performance as CEOs, there was no reason for them to pay for management postgraduate studies.[43] Furthermore, according to the same survey, the forgone earnings for doctors working as CEOs are high, considering their alternative is to work as clinical doctors. The high opportunity cost further disincentivizes doctors to invest in postgraduate management education in the absence of future monetary returns.

[40] Panel D in Figure 2.6 shows that the reform did not impact hospital employees' wages.

[41] The finding that prohibiting individuals with non-medical degrees from becoming CEOs hinders organizational performance is consistent with recent research showing that discrimination against qualified managers can reduce organizational performance (Huber, Lindenthal, and Waldinger, 2021) and, more broadly, that talent misallocation reduces aggregate output. (Hsieh, Hurst, Jones, and Klenow, 2019).

[42] The norm could sustain because CEOs were elected by the head of the Health Service where hospitals are located, who in turn also were doctors and shared the belief that doctors would overperform professional managers.

[43] This is similar to findings in Bloom et al. (2015b), who show that one of the major initial barriers to adoption of management practices was that firms thought they would not be profitable to adopt.

2.6 Conclusion

In this paper, we study the extent to which CEOs in the public sector can improve their organization's performance. We first document that the identity of CEOs matters for public hospital performance in Chile and explains a substantial share of the variation in mortality across hospitals. We then leverage the gradual adoption of a reform which introduced a competitive recruitment process for hiring public sector CEOs, and find that it reduced hospital mortality by approximately 8 percent. We show that this result is not explained by patient selection and is robust to other explanations. In contrast, we find evidence that the reform operates through more efficient use of medical resources and better personnel practices.

We then examine whether the financial incentives included in the reform in the form of performance pay and higher base wages explain the findings. We document that the performance pay incentives are poorly designed and are not binding across the board. Leveraging a later amendment to the reform, we also show that higher wages do not impact managerial performance in our setting. We thus rule out that, conditional on the characteristics of a given CEO, the financial incentives in the reform drive the results.

Instead, we show that the reform displaced older doctors in favor of younger CEOs with educational training in management and that it incentivized doctors who wanted to pursue careers as hospital CEOs to invest in management education. Furthermore, we find that management training is the main predictor of treatment effect heterogeneity. Since this result may be due to differential selection, we view this piece of evidence as suggestive and consider that examining the causal effect of management education on CEO performance in the public sector is an important open question for future research.

To conclude, we note that the reform shifted two different margins of personnel selection that could account for the results. First, conditional on the same pool of individuals willing to take the position, the removal of discretionary appointments in cases in which "outsiders" are implicitly banned from certain positions—which we show was the case in our setting for individuals without medical degrees—is likely to improve the allocation of talent. Second, as discussed above, the extra pay likely plays a role by attracting higher-quality candidates to the pool of applicants. Disentangling these two margins is also a promising avenue for future research.

Figures

Figure 2.1: Adoption of the recruitment process in positions across government agencies

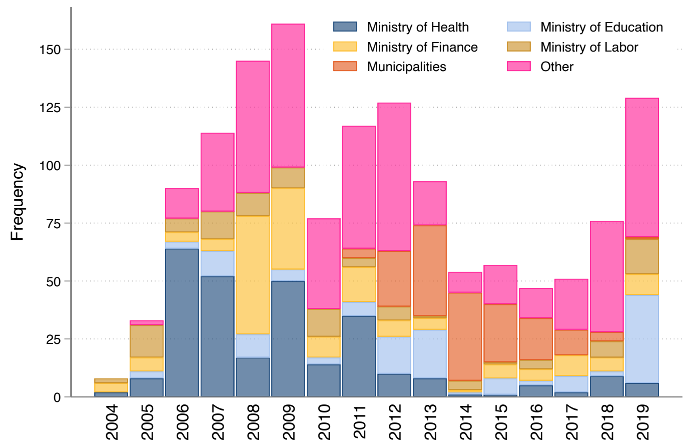

(a) Positions adopting the selection process for first time

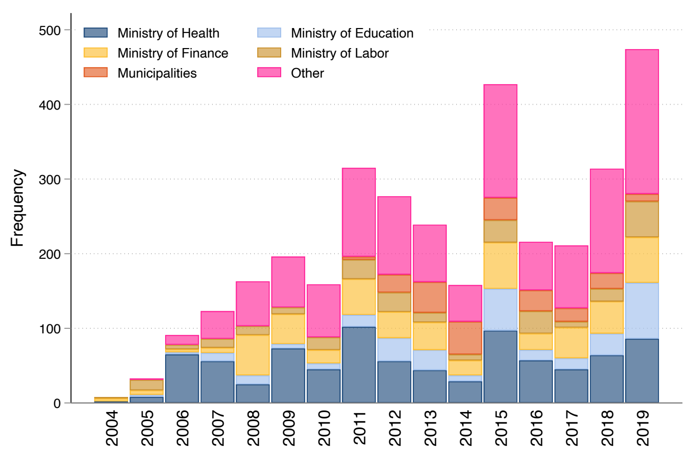

(b) Yearly recruitment processes overseen by the Civil Service

Notes: This figure displays the rollout of the selection reform across government agencies. Panel A presents the number of management positions that adopt the selection reform for the first time. After that, every new manager in that position has to be selected using this mechanism. All senior executive positions created after 2003 have to use the new selection system, and existing positions adopt it gradually. Panel B presents the number of selection processes the Civil Service oversees every year. We use the ending date of the process to allocate the process to a given year. Yearly observations include positions using the selection system for the first time and positions that had already adopted it in the past and are selecting a new manager. The spikes observed in 2011, 2015, and 2019 are evidence of substantial senior executive transitions after a new government is in place.

Figure 2.2: Dynamic effects of the reform on hospital quality

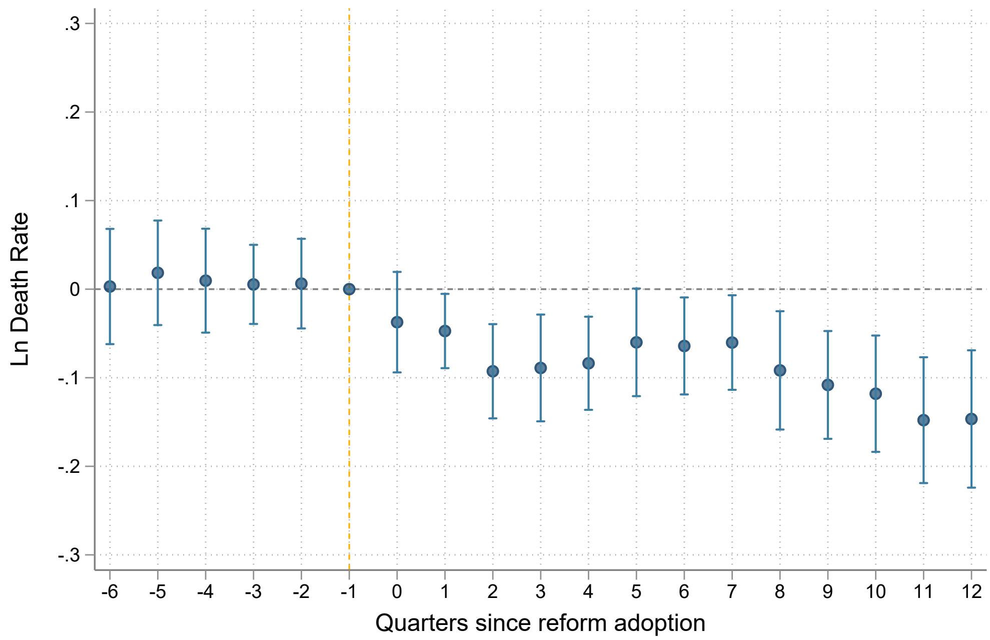

Notes: This figure presents event study evidence of the reform's effect on hospital deaths, following Equation 2.2. The empirical analysis uses quarterly panel data on public hospitals in a time window comprehending 6 quarters before and 12 quarters after the reform was adopted by each hospital, and exploits the gradual adoption of the selection reform in public hospitals during that period. We do not impose a time window for hospitals that did not adopt the policy. Each dot corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. The dashed yellow line represents the omitted coefficient. Standard errors are clustered at hospital level.

Figure 2.3: Testing for patient selection: supply-side

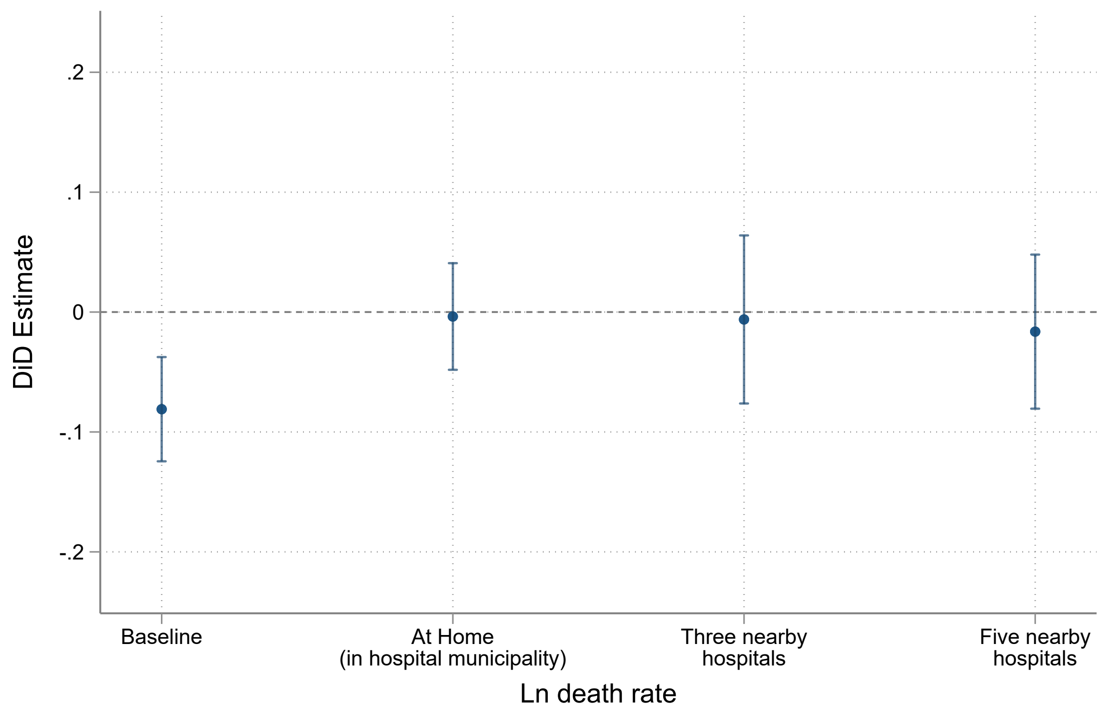

Notes: This figure presents evidence to assess patients' selection as a confounder of our main results. We plot the estimates and confidence intervals obtained by estimating Equation 2.1 for the logged at-home death rate and for logged death rates at nearby hospitals. All regressions consider standard errors clustered at hospital level.

Figure 2.4: Testing for patient selection: demand-side

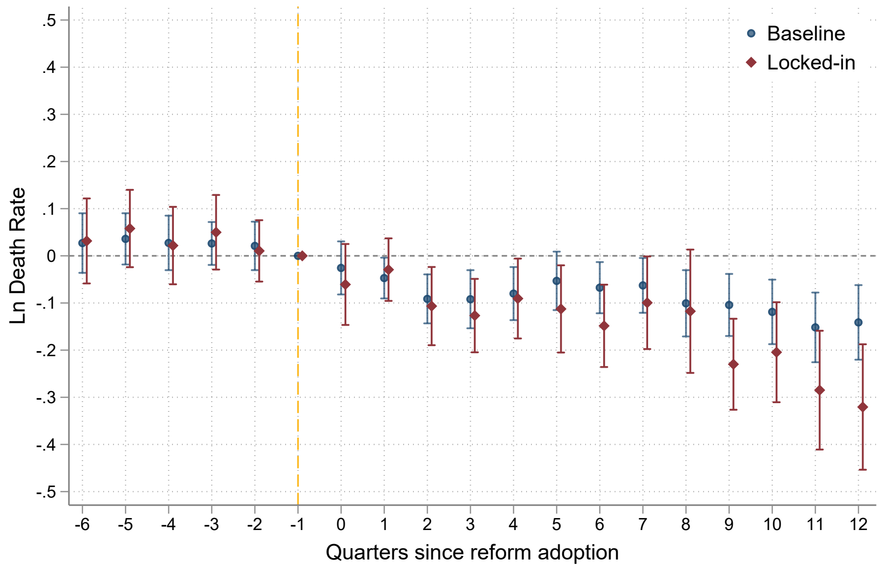

(a) Locked-in patients

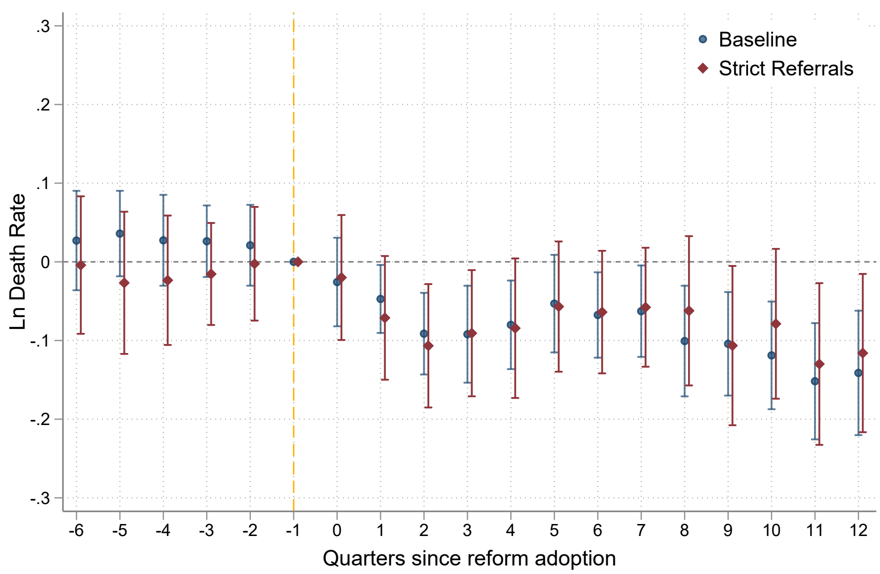

(b) Strict-referrals patients

Notes: This figure presents evidence to assess patients' selection as a confounder of our main results. Panel A presents event study evidence on the reform's effect on hospital deaths, following Equation 2.2, but on a restricted sample of locked-in patients only. Panel B presents event study evidence on the reform's effect on hospital deaths, following Equation 2.2, but on a restricted sample of patients that followed the referrals mandated by the health system. These figures also include baseline estimates for a comparison. All regressions consider standard errors clustered at hospital level.

Figure 2.5: Effect of CEO transition on death rates

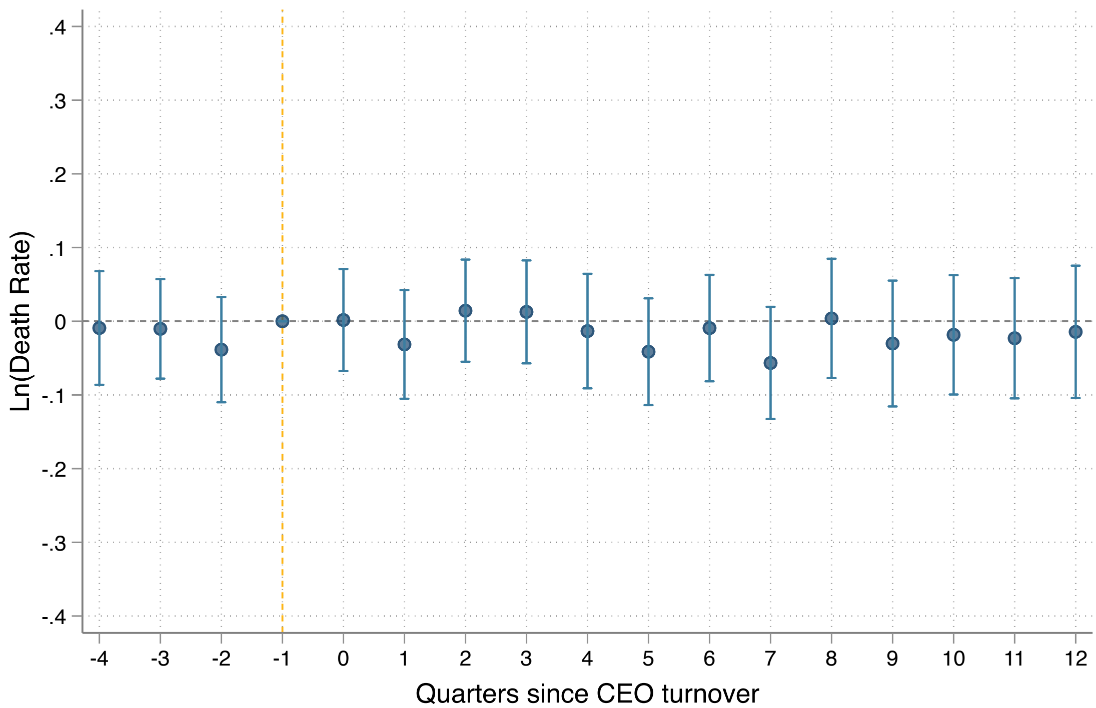

Notes: This figure presents the coefficients of the stacked event study specification in Equation 2.3. An event is a CEO transition in a hospital that never adopts the new selection system. For each transition event, we define a time window around it and a control group of hospitals with no transitions in the time window. We define a set of valid events as those that are balanced in the time window and do not overlap with another transition in the pre-period within the time window. In total, there are 415 valid CEO transitions. The dependent variable is the death rate at the hospital level in a given quarter. The regression includes case mix controls. Dots indicate estimated coefficients and vertical lines the corresponding 95% confidence intervals. Standard errors are clustered at hospital level.

Figure 2.6: How is the reform improving hospital performance?

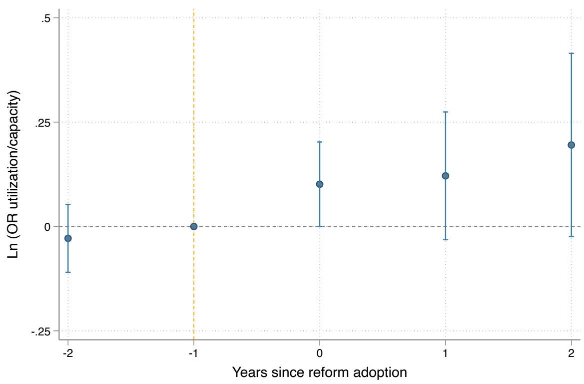

(a) OR utilization

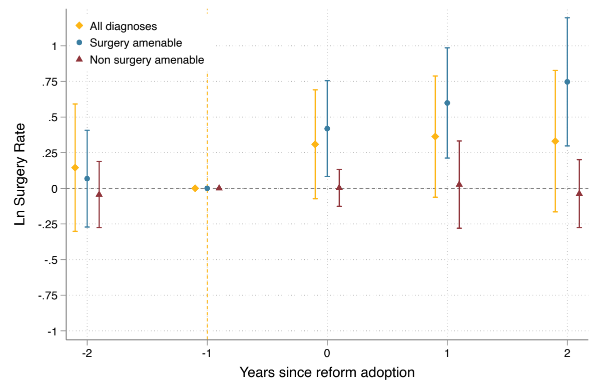

(b) Surgeries

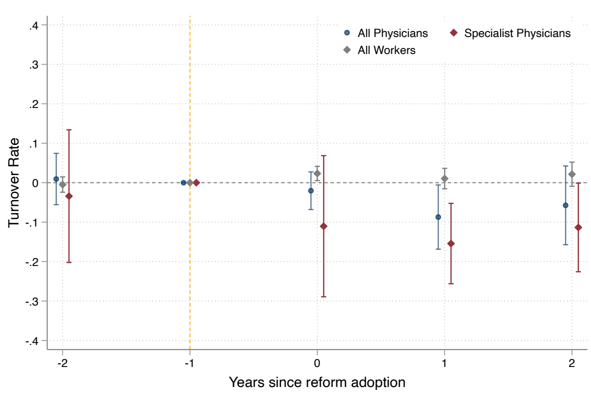

(c) Turnover

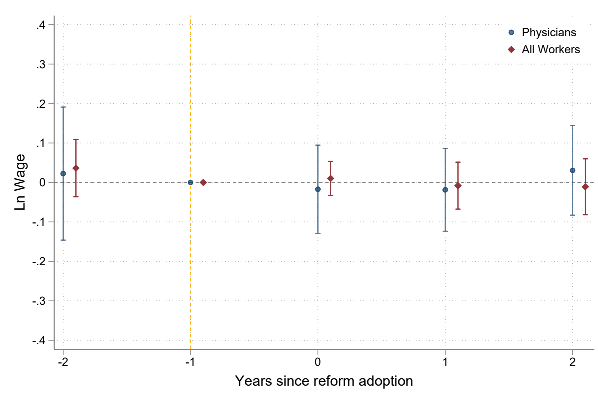

(d) Wages

Notes: This figure presents event study evidence on the reform's effect on several hospital outcomes, following Equation 2.2 estimated at the year level. Panel A examines the logarithm of operation room utilization over capacity. Panel B focuses on logged surgery rates and distinguishes by diagnoses amenable to death prevention through surgery. Panel C replaces the dependent variable with the turnover of doctors (circle markers in blue) and other health workers (diamond markers in red). Turnover is defined as the number of workers in group j who are leaving hospital h in $t+1$ (job to job or job to unemployment transitions) over the number of workers in group j working in h at time t. Panel D consider logged hourly wages as the dependent variable and plots the impact of the policy for wages of all hospital personnel (circle markers in blue) and of doctors (diamond markers in red). Each dot corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure 2.7: The policy displaced doctor CEOs with no management studies

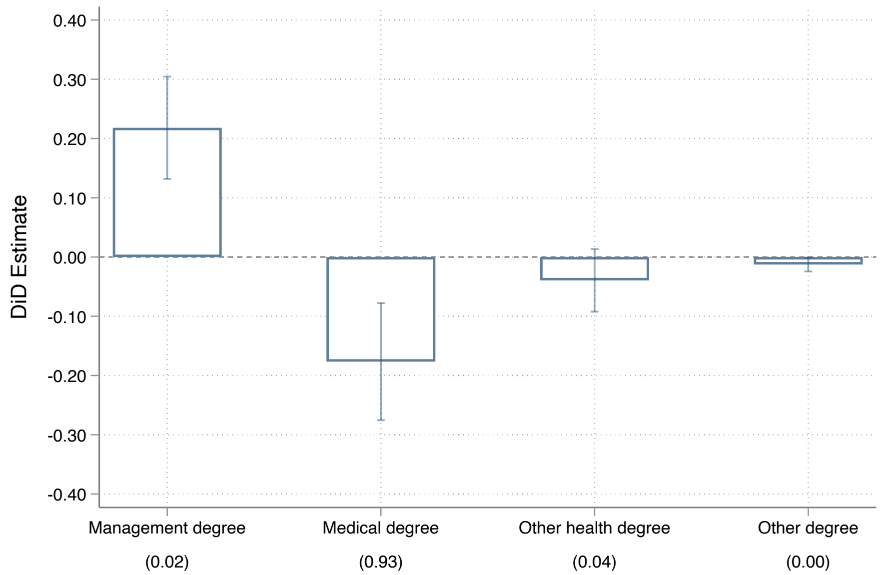

(a) All degrees

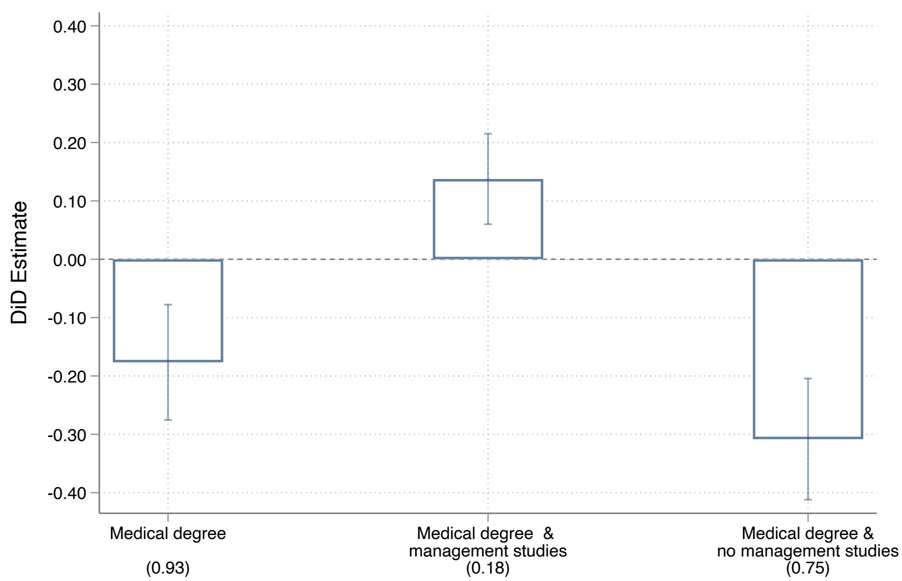

(b) Doctors

Notes: This figure presents the effect of the policy on CEO educational background. Panel A presents the average 3-year effect of the reform on the likelihood that the CEO has an undergraduate management degree, a medical school degree, another health degree, or another major. All categories are mutually exclusive. Panel B focuses on doctors and performs separate estimations for doctors with and without management studies (as of the date of their appointment as CEOs). Bars represent the estimate from Equation 2.1 on each outcome and vertical lines indicate the corresponding 95% confidence intervals. Standard errors are clustered at hospital level.

Figure 2.8: Dynamic effects on CEO educational background

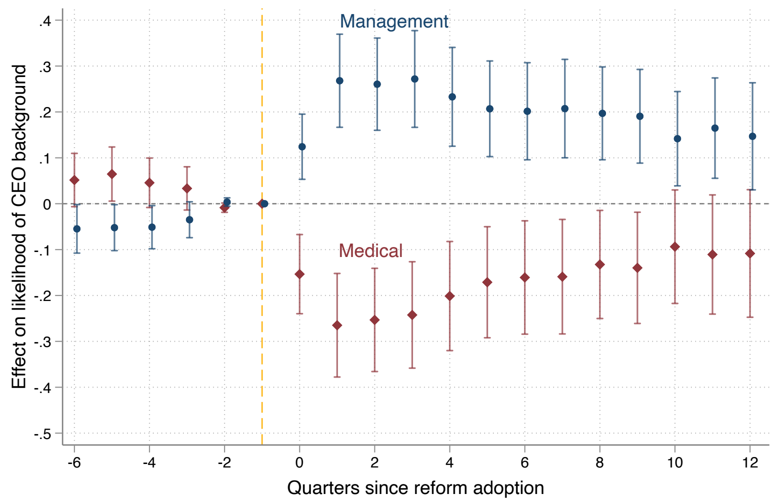

(a) Medical degree vs. management undergraduate degree

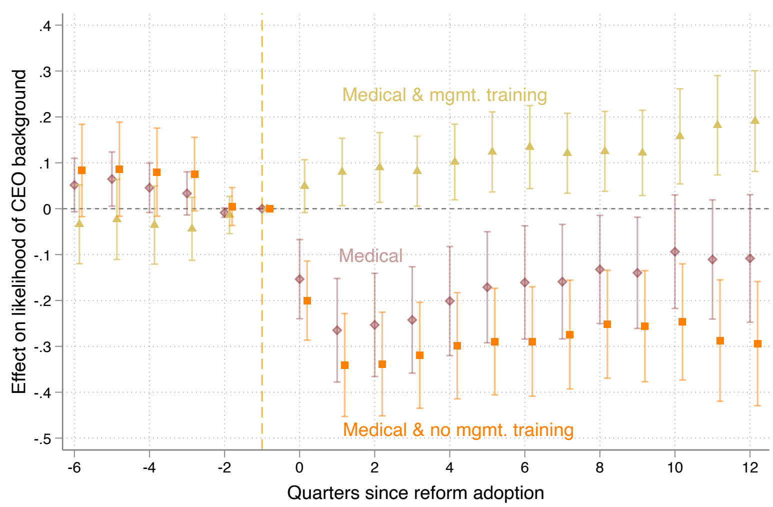

(b) Medical & mgmt. training vs. medical & no mgmt. training

Notes: This figure presents event study evidence of the reform's effect on CEO educational background, following Equation 2.2. In Panel A, the figure overlays the estimation of two dependent variables. The first is a dummy variable that takes value 1 if the CEO has a management-related undergraduate degree (in blue with dot markers). The second corresponds to a dummy variable that takes value 1 if the CEO has a medical degree (in red with diamond markers). Panel B decomposes the total effect on doctor CEOs (in light red with diamond markers) into the change coming from doctor CEOs with management training (the beige triangle markers) and doctor CEOs with no management training (the orange square markers). Management training refers to whether the CEO holds a master's degree or a diploma in management (as of the date of their appointment as CEO). Dots indicate estimated coefficients. The vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure 2.9: The reform made more likely that public hospital CEOs have management training

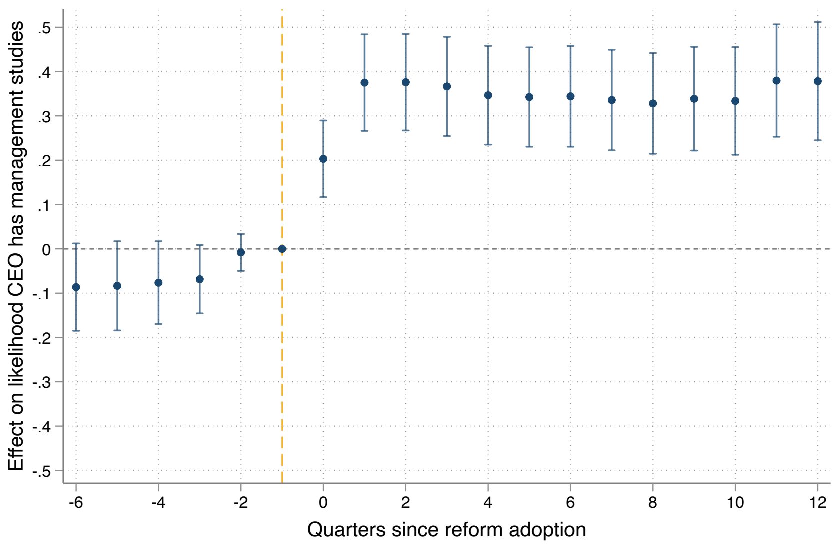

Notes: This figure presents event study evidence of the reform's effect on the likelihood that the CEO has management studies, following Equation 2.2. Management studies is a dummy that takes value 1 if, as of the date of their appointment as CEO, the individual holds an undergraduate degree in a management-related major, or a master's or diploma in management. Dots indicate estimated coefficients. The vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure 2.10: Distribution of performance scores for post-reform CEOs

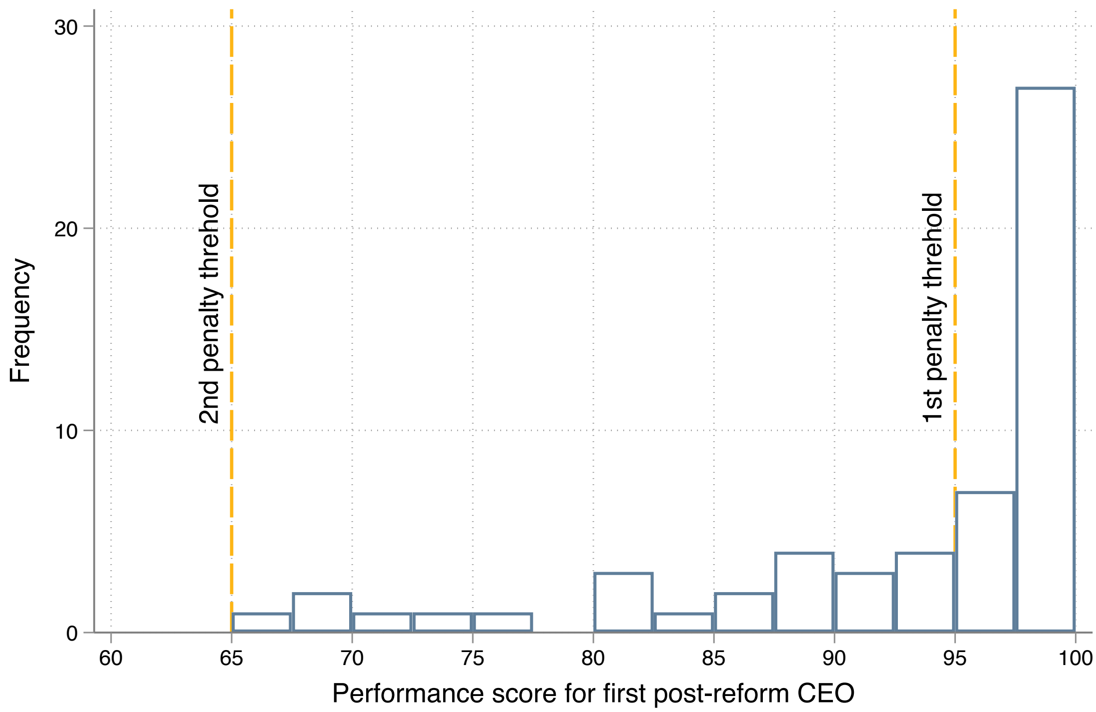

Notes: This figure displays average performance scores for the first post-reform CEO. Before the reform, managerial performance did not affect the wage schedule. After the reform, CEOs face wage penalties if they perform below specific performance thresholds. Performance scores are computed following a 3-year performance contract the CEO defines with her superior. We accessed all available performance contracts and yearly performance scores. Unfortunately, some of the oldest contracts and performance scores are lost, and the Civil Service has no available records. Of the 87 CEOs hired for the first time under the new selection system, we have performance scores for at least 1 year for 57 CEOs. An observation is the average of all available scores for a CEO in her 3-year contract. Dashed yellow lines represent the wage penalty thresholds described in Equation 2.4. Managers who score below the first penalty threshold had to pay a penalty equal to 1.5% of their annual wage. Below the second threshold, the penalty is 7% of their yearly wage.

Figure 2.11: Do efficiency wages impact death rates?

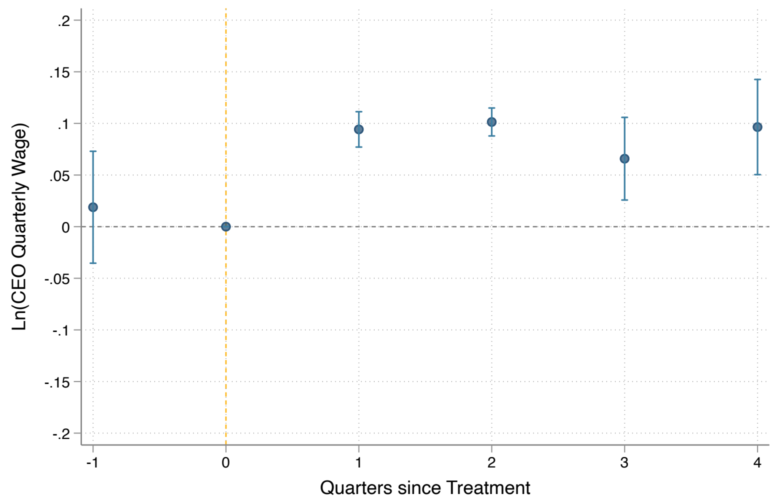

(a) Effect on CEO wages

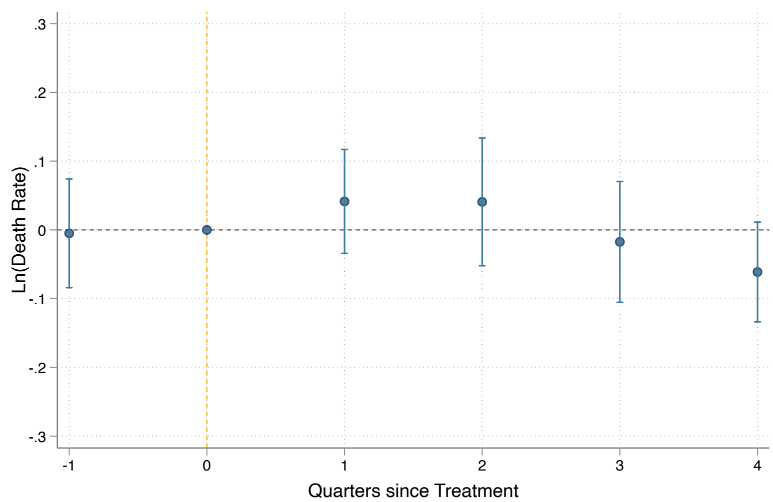

(b) Effect on death rates

Notes: This figure examines the impact of higher hospital CEO wages on hospital performance. The empirical design exploits an amendment to the recruitment reform, which increased wages for CEOs *only if* they are doctors *and* were appointed using the selection reform after November 2016. For each event, we define a time window around the transition and determine an event-specific control group that includes hospitals with no transition and units with transitions to professionals other than doctors. We select valid events that are balanced in the time window and do not overlap with other transitions one period before the event. There are a total of 24 valid events. We then append the data for all valid events and estimate an event study following Equation 2.3. Panel A presents the estimates of the amendment's effect on CEO wages, and Panel B displays the impacts on death rates. The regression on death rates includes include case mix controls. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95% confidence intervals. In Panel A, we cluster standard errors at the CEO's professional degree, which is the treatment unit. In Panel B, we cluster standard errors at hospital level.

Tables

Table 2.1: Explanatory power of managerial talent to account for hospital performance

	Ln Death Rate					
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	6,712	6,712	6,712	6,712	6,712	6,712
R^2	.41	.42	.67	.76	.73	.76
Adj. R^2	.40	.41	.66	.73	.69	.72
Case Mix Controls	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes	Yes	Yes
Hospital FE	No	No	Yes	Yes	No	No
CEO FE	No	No	No	Yes	Yes	No
CEO-by-hospital FE	No	No	No	No	No	Yes
F-statistic for CEO FEs	-	-	-	3.4	10.06	-

Notes: This table shows the extent to which variation in hospital quality can be explained by managerial talent. Panel A compares the adjusted R^2 estimated from several regressions of the logarithm of death rates on different sets of explanatory variables.

Table 2.2: Descriptive statistics

	Mean	Std. Dev.	Bottom 10%	Median	Top 10%	# of Obs.
	(1)	(2)	(3)	(4)	(5)	(6)
Patient Characteristics:						
% Female	0.59	0.08	0.47	0.60	0.68	13,988
% Age < 29	0.36	0.16	0.14	0.37	0.49	13,988
% Age $\in$ (30,29)	0.12	0.05	0.06	0.12	0.17	13,988
% Age $\in$ (40,49)	0.10	0.04	0.06	0.10	0.13	13,988
% Age $\in$ (50,59)	0.10	0.04	0.06	0.09	0.14	13,988
% Age $\in$ (60,69)	0.11	0.05	0.07	0.10	0.16	13,988
% Age $\in$ (70,79)	0.12	0.06	0.06	0.11	0.20	13,988
% Age $\in$ (80,89)	0.09	0.06	0.03	0.07	0.16	13,988
% Age > 89	0.02	0.02	0	0.01	0.05	13,988
% Public Insurance	0.96	0.05	0.92	0.98	1.00	13,988
Hospital Characteristics:						
High-level Hospital	0.33	0.47	0.00	0.00	1.00	13,988
Medium-level Hospital	0.15	0.36	0.00	0.00	1.00	13,988
Low-level Hospital	0.52	0.50	0.00	1.00	1.00	13,988
Total Number of Patients	1,491	2,006	101	587	4,568	13,988
Total Number of Beds	143	177	16	65	415	13,946
Total Number of Surgeries	461	867	0.00	4	1,730	13,988
Physicians per 100 patients	6.75	8.58	2.30	4.91	11.89	6,624
Nurses per 100 patients	6.17	7.72	2.22	4.79	9.89	6,624
Hospital Outcomes:						
Number of Deaths	38.21	63.27	1.00	12.00	116.00	13,988
Death Rate	2.46	1.94	0.38	2.15	4.69	13,988
Death Rate 28 days	4.21	2.87	1.18	3.66	7.83	13,988
Death Rate ER	3.01	3.53	0.15	2.55	5.69	11,087
Death Rate ER AMI	12.21	23.77	0.00	2.38	33.33	4,555

Notes: This table presents descriptive statistics for the universe of public hospitals included in our main analysis. Patient characteristics and hospital outcomes come from individual-level inpatient records collected by the Ministry of Health, and encompass almost 29 million hospital events (DEIS, 2019). Hospital characteristics come from hospital-level public records, and restricted-use administrative data covering the universe of employees in all public hospitals between 2014 and 2019, which is collected by Ministry of Health for HR purposes.

Table 2.3: Balance in observable characteristics before the reform

	Avg. never adopter	β Ever adopter (Levels)	β Ever adopter (First-Diff)
	(1)	(2)	(3)
Patient composition:			
% Age < 29	0.381	0.042	0.004
		(0.060)	(0.003)
% Age $\in$ (30,49)	0.220	0.005	0.003
		(0.021)	(0.002)
% Age $\in$ (50,69)	0.185	0.009	-0.003
		(0.024)	(0.003)
% Age $\in$ (70,89)	0.197	-0.047**	-0.004*
		(0.021)	(0.002)
% Age > 89	0.018	-0.009***	-0.000
		(0.002)	(0.001)
% Female	0.605	-0.027	0.000
		(0.018)	(0.003)
% Public insurance	0.972	-0.043***	0.003
		(0.009)	(0.002)
Hospital outcomes:			
Number of deaths	5.970	47.943***	0.999
		(16.157)	(1.053)
Death rate	1.389	0.497	0.083
		(0.366)	(0.083)
Death rate ER	1.483	1.325**	0.137
		(0.618)	(0.116)
Death rate 28 days	3.305	-0.046	0.155
		(0.504)	(0.143)
Death rate AMI	23.465	-19.993***	6.085
		(6.446)	(14.883)
Political variables:			
% Votes for right	25.764	8.186*	2.674
		(4.792)	(5.691)
% Votes for center	19.107	5.499	2.046
		(5.633)	(3.970)
% Votes for left	24.435	-8.226	-4.579
		(5.256)	(4.275)

Notes: This table studies differences between ever- and never-adopter hospitals in terms of predetermined characteristics. We consider a window of six quarters before adoption. Column (1) shows the average of each characteristic for never adopters. Column (2) presents the coefficient obtained from a regression of each variable on a dummy that equals 1 if the hospital was an ever adopter. Column (3) replicates column (2) but replaces the dependent variable with its first differences. The political variables correspond to the vote share of right wing, center, and left wing parties in the 2000 and 2004 mayoral elections in the municipalities where hospitals are located. The first differences of these variable correspond to the difference in vote shares between the 2000 and 2004 elections. Standard errors are clustered at the hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.4: Impact of the reform on death rates

	Ln Death Rate					Poisson (# Deaths)	
	All			28-days	ER	All	ER: AMI
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1 if reform adopted	-0.131***	-0.081***	-0.095***	-0.061***	-0.156***	-0.054***	-0.146
	(0.025)	(0.022)	(0.023)	(0.016)	(0.036)	(0.018)	(0.134)
Observations	8,104	8,104	8,104	8,104	6,592	8,104	1,956
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hospital FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Case-Mix Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Flexible Case-Mix Interactions	No	No	Yes	Yes	Yes	Yes	No
# of Hospitals	181	181	181	181	175	181	132
Mean Dep. Variable	2.625	2.625	2.625	4.726	3.088	21.85	16.22

Notes: This table presents our estimates of the impact of the selection reform on public hospital's performance, as measured by death outcomes. Estimates are from the staggered DiD specification in Equation 2.1. The empirical analysis uses quarterly panel data for public hospitals in a time window comprehending 6 quarters before and 12 quarters after the reform was adopted by each hospital, and exploits the gradual adoption of the selection reform in public hospitals during that period. We do not impose a time window for hospitals that did not adopt the policy. In columns (1)-(3), we focus on in-hospital death rates and add case mix controls sequentially. Column (4) replaces the dependent variable with 28 days after admission death rate, and thus considers in- and out-of-hospital deaths. In column (5) we study the impact of the reform on death rates of ER admissions. Finally, columns (6) and (7) reports estimates from a Poisson regression of death counts. Column (7) focuses on the subset of emergency room admissions with AMI (Acute Myocardial Infarctions, commonly known as "heart attacks") diagnoses. Results in columns (1)-(6) are weighted by the number of the hospital's inpatients as of 2005. For columns (1)-(6), the mean dependent variable is presented in levels instead of logs. Standard errors are displayed in parentheses and are clustered at hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.5: CEO selection reform v. other policies

Policy (1)	Paper (2)	Death rate definition (3)	Average death rate (4)	Impact on death rate (5)	Sample of patients (6)
Spending					
↑ 10% p/capita	Doyle et al. *JPE* '15	All, 1-year	37%	↓ 6%	ER + Amb. + ≥ 65*
	Ours		32%	↓ 7%	ER + ≥ 65
Public vs Private					
VA v. Non-VA hospitals	Card & Chan '22	All, 1-year	29%	↓ 7%	ER + Amb.+ ≥ 65
	Ours		32%	↓ 7%	ER + ≥ 65
Competition					
+1 hospital in neighborhood	Bloom et al. *ReStud* '15	In-hospital, 28-day	15%	↓ 10%	ER + AMI
↓ 10% HHI	Gaynor et al. *AEJ EP* '13	In-hospital, 28-day	1.6%	↓ 1%	All patients
	Ours		2.3%	↓ 15%	All patients

Notes: This table compares the impact of the CEO selection reform we study with the impact of other policies previously studied in the literature. To construct this table, we estimate our main Equation 2.1 for the different dependent variables—reported in column (3)—and in different samples of patients reported in column (6). For more details, see Subsection 2.3. Acronyms used in the table: ER: Emergency Room; AMI: Acute Myocardial Infarction; Amb: arriving by ambulance; *: non-deferrable medical condition.

Table 2.6: Effect of the reform on managers' skills and demographics

	Skills				Demographics	
	CEO Fixed Effect	Avg. Test Score	Math Specific Exam	Science Specific Exam	Age	Female
	(1)	(2)	(3)	(4)	(5)	(6)
1 if reform adopted	-0.09***	-0.12	0.08	-0.13**	-1.87*	-0.03
	(0.03)	(0.10)	(0.08)	(0.05)	(1.06)	(0.05)
Observations	4,391	7,053	5,561	5,561	7,906	8,085
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Hospital FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	No	No
# of Hospitals	111	177	162	162	180	180
Mean Dep. Variable	0.570	2.000	0.740	0.990	50.190	0.210

Notes: This table presents our estimates of the impact of the selection reform on public hospital CEOs' skills and demographics. Estimates are from the staggered difference-in-differences specification in Equation 2.1, but we switch the dependent variable for CEO characteristics. The empirical analysis uses quarterly panel data between 2001 and 2019 and exploits the gradual adoption of the selection reform by public hospitals in that period. In columns (1), we focus on our CEO fixed effects estimates as a measure of managerial ability. Columns (2)-(4) examine the impact on college admission test scores as a proxy for cognitive abilities. The math- (science-) specific exam takes value 1 if the manager took the math- (science-) specific exam in the older version of the college entrance exam in Chile, in which applicants had to choose which exam to take. Columns (5)-(6) study the effect on the age and gender of the CEO. The mean dependent variable is computed in the period before each hospital adopted the reform. Standard errors are displayed in parentheses and are clustered at hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.7: Heterogeneity in CEO performance by managerial education background

Identifying variation is:	Reform adoption			CEO transition	
	Ln Death (%)	Ln Death (%)	Ln Death (%)	Ln Death (%)	Ln Death (%)
	(1)	(2)	(3)	(4)	(5)
Reform & mgmt. undergrad.	-0.111***				
	(0.029)				
Reform & non-mgmt. undergrad.	-0.076***				
	(0.026)				
Reform & any mgmt. studies		-0.122***	-0.130***		
		(0.025)	(0.028)		
Reform & non-mgmt. studies		-0.028	-0.027		
		(0.027)	(0.027)		
CEO with management studies				-0.072***	
				(0.025)	
CEO with no management studies					-0.010
					(0.022)
Sample	All CEOs	All CEOs	Doctor CEOs		
Observations	8,085	8,085	5,732	71,027	193,177
Time FE	Yes	Yes	Yes	Yes	Yes
Hospital FE	Yes	Yes	Yes	Yes	Yes
Case mix Controls	Yes	Yes	Yes	Yes	Yes
# of Hospitals	181	181	176	168	175
Mean Dep. Variable	2.63	2.63	2.49	2.88	2.41
p-value *Mgmt. = Non Mgmt.*	0.22	0.00	0.00		

Notes: This table examines heterogeneous effects of the reform by CEOs managerial education background. Columns (1)-(3) focus on the differential effect of the selection reform on death rates, following the staggered DiD design in Equation 2.1. In Panel A, we ask to what extent the reform has differential effects depending on the CEO's educational background. Column (1) interacts adoption of the selection reform with whether the CEO holds an undergraduate degree in a management-related undergraduate major. Columns (2) and (3) focus on whether the CEO has *any* management studies, which include undergraduate and postgraduate studies related to management. Columns (4)-(5) present the results of the stacked event study specification in Equation 2.3. In column (4), an event is a transition from a CEO without management studies to a CEO with management studies. In column (5), an event is a transition from a CEO without management studies to a CEO without management studies. For each transition event, we define a time window around it and a control group of hospitals with no transitions in the time window. We define a set of valid events as those that are balanced in the time window and do not overlap with another transition in the pre-period within the time window. We also exclude transitions associated with the first time that a CEO was appointed after the selection reform was adopted by a given hospital. In total, there are 94 valid CEO transitions, as described in Appendix Table B.6.4. The dependent variable is the death rate at hospital level in a given quarter. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95% confidence intervals. Standard errors are clustered at hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Chapter 3

Equilibrium Effects of Food Labeling Policies

3.1 Introduction

Obesity rates in the world have tripled over the last half-century. Today, about 40% of the world's adult population is either obese or overweight (WHO, 2018). One increasingly popular policy tool governments are using to combat obesity are front-of-package labels, which are visual warnings placed prominently on the front of packaged food products. Unlike nutrition facts tables, which provide detailed information on the back of food products, food labels are simple symbols that clearly signal to consumers when a particular product is considered unhealthy. Since 2016, more than 25 countries have either implemented or are in the process of implementing country-wide mandatory food labeling policies (Barahona, Kim, Otero, and Otero, 2022).

Several features of food labels make them popular. First, providing information to consumers is widely perceived as innocuous, in the sense that it can only improve consumer welfare. Furthermore, sugar taxes—the most prominent instrument to combat obesity—may be regressive (Allcott, Lockwood, and Taubinsky, 2019c). Finally, in settings in which some but not all agents act against their own interest, information interventions can be more efficient than taxes because their effects are better targeted (Bernheim and Taubinsky, 2018). Opponents of food labels, however, argue that they are ineffective in improving consumers' diet and impose an unnecessary burden on firms.

Most of this discussion focuses on consumers' responses to labels. However, firms' responses to the large-scale implementation of food labels may undo or even amplify some of their desirable properties. Food labels can, for example, affect product differentiation and market power. Firms may also use healthier ingredients in their products to avoid receiving labels, thus amplifying the positive effects on nutritional intake but also increasing consumer prices as a result of increased production costs. Taken together, the impact of large-scale food labeling regulations is ambiguous.

This paper studies the equilibrium impacts of food labels on consumers' purchases, firms' pricing and production decisions, nutritional intake, and consumer welfare. We combine descriptive analyses with a model of supply and demand for food and nutrients to quantify the impact of the Chilean Food Act of 2016, the first mandatory nationwide food labeling regulation implemented in the world. The regulation mandates that food manufacturers put warning labels on all of their packaged food products that surpass a threshold concentration of sugar, calories, sodium, or saturated fat.

To study how the regulation affected consumer choice, we use scanner data on purchases made in Walmart, the largest food retailer in Chile, from 2015 to 2018. The data contain information on prices, quantities, and consumer demographics such as gender, age, and income. To shed light on mechanisms, we surveyed 1,500 consumers and elicited their beliefs over the nutritional content of products. Finally, we use scanned nutrition facts tables of products before and after the policy to study strategic reformulation decisions by firms. We thus have a rich window into consumer demand and beliefs, as well as firm behavior.

We focus our analysis on the breakfast cereal market. Cereal is well suited for this analysis because it is a well-defined category with little substitution across other food categories, substantial labeling variation across products, and one in which food labels may be particularly informative due to consumers' nutritional content misperceptions. We extend the analysis to other product categories in Barahona et al. (2022).

Three key findings arise from our descriptive analysis. First, we show that consumers substituted from labeled to unlabeled products. Second, we find that the change in demand is primarily driven by updates in consumer beliefs. Products that consumers already knew had high sugar or caloric concentration only experienced a small and temporary drop in demand. However, products that consumers previously believed to be low in sugar and calories but received a label under the labeling policy experienced a persistent 40% decrease in demand relative to unlabeled products. In line with a Bayesian updating model, this result suggests that labels are more effective when they provide new information to consumers. Third, we find that suppliers responded to the regulation by reformulating their products and changing prices. To avoid labels, many firms modified the nutritional content of their products to be just below the regulatory thresholds and decreased sugar and caloric concentration by 11.5% and 2.8%, respectively. We also document a 5.5% increase in prices of unlabeled products relative to labeled ones due to the regulation.

Motivated by these findings, we develop and estimate a model of supply and demand for food and nutrients. On the demand side, consumers care about the price, taste, and healthiness of products. Healthiness, however, is not observed, and consumers may have poorly calibrated beliefs about products' nutritional content. Food labels help consumers by providing them with a binary signal about the true nutritional content of products, which allows them to make better-informed purchasing decisions. On the supply side, firms strategically set prices and nutritional content to maximize profits. Food labels create a sharp discontinuity in demand at the policy threshold, which induces firms to reformulate their products to avoid labels. However, reducing the concentration of critical nutrients is costly, and may cause firms to raise prices.

Our model highlights two sources of inefficiency that arise due to incomplete information. First, consumers may make mistakes when choosing what to buy. Second, firms do not have incentives to produce healthier items if they cannot credibly inform consumers about product healthiness. Thus, food labels may reduce inefficiencies by improving consumer choice and incentivizing suppliers to produce healthier goods.

We use our model to quantify the impact of the Chilean Food Act on nutritional intake and consumer welfare. To analyze how equilibrium forces change the effectiveness of food labeling policies, we simulate three progressively more flexible counterfactuals, each of which we benchmark against a no-intervention counterfactual.

First, we study the effects of food labels in the absence of supply-side responses. We find that the regulation reduces sugar and caloric intake in the cereal market by 6.8% and 0.6%, respectively, resulting in average gains in consumer welfare equivalent to 1.1% of total cereal expenditure. The changes in consumer welfare are driven by a combination of a healthier diet, fewer dollars spent, and an increase in the consumption of less tasty products (e.g., oatmeal).

Second, we allow firms to optimally set prices in response to the policy but not to change the nutritional content. As in Villas-Boas, Kiesel, Berning, Chouinard, and McCluskey (2020), we use this counterfactual to assess the role of product differentiation and market power. Under this counterfactual, prices of unlabeled and labeled products go up and down, respectively, with average prices remaining relatively constant. Gains in consumer welfare relative to the no-intervention counterfactual are 7% lower than in the absence of supply-side responses.

Third, we allow firms to optimally reformulate their products to avoid receiving labels. This counterfactual recovers the full effect of the policy. Overall, we find that high-in-taste products become healthier but more expensive due to higher production costs. Consumer welfare gains under this counterfactual are 70% larger than in the absence of supply-side responses.

We then use our model to study optimal policy design. We show that ignoring supply-side effects can lead to substantially different outcomes. Considering only demand-side effects, a social planner who wants to maximize consumer welfare should set a threshold that maximizes the information provided by labels. However, when accounting for supply-side responses, the social planner wants to set a lower threshold to provide stronger incentives for firms to improve the nutritional content of their products. By taking supply-side responses into account, the social planner can reduce sugar intake by an additional 38% and increase consumer welfare gains by 20% relative to the outcome under the threshold that maximizes information.

Overall, our descriptive and model results suggest that food labels are more effective when consumers have mistaken beliefs about products' healthiness, consumers value healthiness, reformulation that does not substantially change products' taste is feasible, and regulatory thresholds are set so that they provide useful information to consumers and encourage product reformulation.

Finally, we compare food labels with other popular policy instruments, such as sugar taxes. When compared with sugar taxes, food labels present both advantages and disadvantages. They tend to be more progressive and better targeted, but are less effective against non-informational market imperfections, such as lack of self-control or fiscal externalities.

This paper contributes to several strands of the literature. It adds to a large literature that studies consumer choice in settings of imperfect information (Hastings and Weinstein, 2008; Abaluck and Gruber, 2011; Abaluck, 2011; Woodward and Hall, 2012; Handel and Kolstad, 2015; Allcott and Knittel, 2019). Moreover, it contributes to the literature that examines how providing nutritional information affects consumer demand. This includes consideration of the effects of advertising (Ippolito and Mathios, 1990, 1995; Dubois, Griffith, and O'Connell, 2017); nutritional information on menus (Wisdom, Downs, and Loewenstein, 2010; Bollinger, Leslie, and Sorensen, 2011; Finkelstein, Strombotne, Chan, and Krieger, 2011); and food labeling regulations (Kiesel and Villas-Boas, 2013; Zhu, Lopez, and Liu, 2015; Allais, Etilé, and Lecocq, 2015). Previous research has also highlighted the importance of firms' strategic responses to nutritional information policies by adjusting prices (Villas-Boas et al., 2020) and reformulating products (Moorman, Ferraro, and Huber, 2012; Lim, Rishika, Janakiraman, and Kannan, 2020). Our paper contributes to these studies by providing evidence of and quantifying the equilibrium effects of national information policies, by allowing firms to vary prices and nutritional characteristics of the products they sell.

Other concurrent work has also studied the Chilean Food Act. Using a before-after analysis, Taillie, Reyes, Colchero, Popkin, and Corvalán (2020) document a significant decline in purchases of labeled beverages following the policy's implementation. Araya, Elberg, Noton, and Schwartz (2022) take advantage of the staggered introduction of labeled products in store inventories and find that labels decrease demand in the breakfast cereal category, but not for chocolates or cookies. Pachali, Kotschedoff, van Lin, Bronnenberg, and van Herpen (2022) study price adjustments and conclude that prices of labeled products increased due to increased product differentiation. Alé-Chilet and Moshary (2022) provide evidence of bunching just below regulatory thresholds and conclude that reformulation reinforces the policy's effects by lowering the caloric content of cereal. Our paper goes further along several dimensions. First, we develop an equilibrium framework that allows both price adjustments and product reformulation. This is crucial in assessing the overall role of equilibrium responses to food labeling policies. Second, we show that beliefs over nutritional content are a primary driver of consumer behavior and explicitly incorporate them in our model. This allows us to provide a welfare evaluation of the policy. Third, we use our model to answer additional policy-relevant questions, such as the design of optimal policy thresholds and the comparison of food labels with sugar taxes. Barahona et al. (2022) combine the insights of this paper with analysis of other product categories and discuss the effectiveness of food labeling policies in different settings.

Our work also relates to the literature on quality disclosure and certification that studies the effect of third-party disclosure on consumer choice and seller behavior (Dranove, Kessler, McClellan, and Satterthwaite, 2003; Jin and Leslie, 2003; Greenstone, Oyer, and Vissing-Jorgensen, 2006; Dranove and Jin, 2010; Roe, Teisl, and Deans, 2014; Houde, 2018; Vatter, 2021) and to the literature in industrial organization that estimates demand models under endogenous product characteristics (Ackerberg and Crawford, 2009; Draganska, Mazzeo, and Seim, 2009; Fan, 2013; Wollmann, 2018).

Finally, we contribute to a broader literature that studies how governments can help consumers make better nutritional choices. Allcott, Diamond, Dubé, Handbury, Rahkovsky, and Schnell (2019a) study whether improving access to healthy food in poor neighborhoods can decrease nutritional inequality, Dubois et al. (2017) analyze the effect of advertising on junk food consumption, and several other papers study the effects and design of taxes for sugar-sweetened beverages and calorie-dense food products (Falbe, Rojas, Grummon, and Madsen, 2015; Falbe, Thompson, Becker, Rojas, McCulloch, and Madsen, 2016; Silver, Ng, Ryan-Ibarra, Taillie, Induni, Miles, Poti, and Popkin, 2017; Allcott et al., 2019c; Lee, Falbe, Schillinger, Basu, McCulloch, and Madsen, 2019; Taylor, Kaplan, Villas-Boas, and Jung, 2019; Dubois, Griffith, and O'Connell, 2020; Aguilar, Gutierrez, and Seira, 2021). Our paper focuses on a different policy instrument and shows that it can be an effective tool to improve diet quality and combat obesity.

The remainder of the paper is organized as follows. Section 3.2 describes the setting and the data. In Section 3.3, we provide descriptive evidence to illustrate the main mechanisms through which food labels can reduce the intake of critical nutrients. In Sections 3.4 and 3.5, we present and estimate the demand and supply model, respectively. We present our main counterfactual exercises in Section 3.6 and conclude in Section 3.7.

3.2 Setting and Data

The Chilean Food Act

In 2015 the Chilean legislature, concerned about the growing obesity problem, passed Law 20.606 (hereafter, the Food Act) to improve nutritional choices. The Act imposed new regulations on how food manufacturers could package and advertise food products. An important part of the Act was a food labeling system, which prominently informs to consumers which products are considered unhealthy.[1] The Food Act sought to enhance consumers' decision-making by providing easy-to-process information about the healthiness of food products.

The Food Act established threshold values for sugar, calories, sodium, and saturated fat concentration and mandated suppliers to place a warning label on the front of their packaged products for each nutrient threshold surpassed. The thresholds were implemented in three stages, with each stage setting stricter threshold values than the last. Due to data limitations, we focus on stage 1, which was implemented in June of 2016 and established limits of 22.5 grams of sugar and 350 kcal per 100 grams of product.[2]

[1] The Food Act also included a ban on selling, distributing, or advertising labeled products in schools, and a ban on advertising labeled products aimed at children younger than 14 years old.

[2] The law was first approved in Congress in 2012 and its details were finalized and announced in June of 2015, one year before Stage 1. Stages 2 and 3 took place in June of 2018 and 2019, respectively. The thresholds were established based on the 90th percentile of the distribution of the concentration of critical nutrients

Data

We restrict our attention to breakfast cereal because it is a well-defined category with substantial labeling variation; around 60% of cereal products received at least one label. Breakfast cereal is also a category in which consumers tend to have inaccurate beliefs about the healthiness of products. This feature is important because, as shown below, beliefs play a critical role in the extent to which labels impact shoppers' decisions. In certain other categories, such as soft drinks, products have already long been categorized as diet and non-diet, and consumer beliefs about nutritional content are thus more closely aligned with reality.[3]

Walmart data

To capture prices and quantities, we use scanner-level data provided by Walmart-Chile. Walmart is the largest food retailer in Chile and accounts for more than 40% of supermarket sales. Our data contain all transactions that occur in any Walmart store in Chile between May 2015 and March 2018. Every transaction identifies products at Universal Product Code (UPC) level and contains information about price, revenue, product name, brand name, and discounts. We can track buyers enrolled in Walmart's loyalty program and link them to individual characteristics, such as gender, age, and household income. We supplement these data with additional information about product and store characteristics also provided by Walmart.

Since our data only cover purchases at Walmart and most consumers may also purchase a large share of their groceries from other retailers, we restrict our analysis to regular Walmart customers. Our final sample consists of 524,000 consumers who visited a Walmart store at least once every 8 weeks during the study period. The average customer in our panel is 48 years old, and 69% are women.[4] In the first year of data, from May of 2015 to May of 2016, the average customer buys cereal 11 times and spends a total of $25 on it.

Nutritional Information

Nutritional data for packaged products come from two sources: (a) pre-policy data collected by the Institute of Nutrition and Food Technology (INTA) at the University of Chile, and (b) post-policy data that we collected and digitized ourselves. The data comprise information on 94 cereal products, which represent 94% of total cereal revenue.

---

from non-processed food products using data from the United States Department of Agriculture (USDA). As far as we know, the choice of thresholds was not influenced by the industry's lobby. The legislation only applies to processed and packaged foods. This means that products that do not have any added sugar, sodium, saturated fat, honey, or syrup do not receive a label, even if they are above a given threshold. For example, even though oats have a caloric content above 350 kcal/100 g, they did not receive a label.

[3] In Barahona et al. (2022), we extend the analysis to several other categories. We also study potential between-category substitution effects and find no evidence of it.

[4] The sample is fairly representative of the Chilean urban population, with high-income consumers slightly overrepresented. A third of consumers are in the bottom 50% of the national income distribution, a third between the 50th and 85th percentiles, and a third in the top 15%.

Consumer beliefs

We conducted a survey to elicit consumers' beliefs about the nutritional characteristics of all cereal products in the absence of food labels. We implemented the survey in Argentina using Qualtrics in August 2019 and surveyed a total of 1,500 individuals. We asked consumers to provide their best estimate of the sugar and caloric concentration of all cereal products and to state how confident they were about their answers. Using this information, we elicit the first and second moments of consumer beliefs about each product's nutritional content. We also collected information about the gender, age, and household income of survey respondents.

We find that, on average, individuals have relatively accurate beliefs about the concentration of sugar in cereal. The correlation between actual sugar content and respondents' stated beliefs is 0.76. However, respondents' beliefs about the caloric concentration of cereal were less aligned with reality; the correlation between the actual and predicted caloric concentration is only 0.26.

3.3 Descriptive Evidence

This section provides descriptive evidence of the impact of the food labeling policy on nutritional intake, consumer choice, and firm behavior. For our analysis, we define a product as the union of UPCs that share the same product name and brand. For example, we assign all *Honey Nut Cheerios* the same product ID regardless of their box size. In total, our sample contains 94 unique cereal products (produced by 14 firms): 39 did not receive a label and 55 received a high-in-calories label, of which 21 received an additional high-in-sugar label. No cereal products received a high-in-sodium or high-in-fat label in our sample period. Our main analysis focuses specifically on caloric and sugar intake. We assign labels to a product based on its 2018 nutritional content.

Three key facts emerge from the evidence presented below. First, consumers decreased demand for labeled products relative to unlabeled ones. Second, products that were perceived as healthy but received labels experienced the largest decline in demand. Third, suppliers responded to the policy by reformulating their products and changing prices.

Changes in equilibrium quantities

We quantify the effects of the policy on demand by using an event-study design. We aggregate our data into product-store-period data bins (where a period is defined as eight consecutive calendar weeks) and estimate the following regression:

$$\log(q_{jst}) = \sum_k \beta_k \cdot L_j \cdot \mathbb{1}\{k = t\} + \gamma \cdot \log(p_{jst}) + \delta_{js} + \delta_t + \varepsilon_{jst}, \tag{3.1}$$

where q_{jst} denotes the grams of product j sold in store s in period t, p_{jst} refers to the product's price per 100 grams of cereal, and L_j is an indicator variable that takes the value of one if the product has one or more labels. Finally, δ_{js} refers to product-store fixed effects and δ_t to period fixed effects. We normalize the β_k coefficients so that their average value over the pre-policy period is equal to zero. Observations are weighted by product-store pre-policy revenues. Products that do not appear in the pre-period have zero weight and are thus excluded from the estimation sample. Standard errors are clustered at the product level.

Figure 3.2(a) displays the results of estimating Equation (3.1). In the pre-period, the coefficients are small and not significantly different from zero. After the regulation was implemented, the quantity of labeled products sold relative to unlabeled ones decreased by an average of 26.4%. The impact of the legislation does not seem to change over time. This suggests that labels shifted consumer purchases away from labeled products, with the effect lasting throughout the entire period covered by our sample.

The role of beliefs

To investigate how information and beliefs shape consumer choices, we use the beliefs survey described in Section 3.2. We use the elicited beliefs about caloric concentration to test for heterogeneity in the impact of labels. If labels provide useful information for consumers, then products for which labels come as a surprise (i.e., products that consumers believed were low in calories but are actually high in calories) should experience a larger drop in demand. We thus split our sample of labeled products into two groups: products below the median in the distribution of beliefs (20 products) and products above the median in the distribution of beliefs (21 products). We use indicator dummies for each of these groups (denoted by Low_j and $High_j$) to estimate the following equation:

$$\log(q_{jst}) = \sum_k (\beta_k^l \cdot L_j \cdot Low_j + \beta_k^h \cdot L_j \cdot High_j) \cdot \mathbb{1}\{k = t\} + \gamma \cdot \log(p_{jst}) + \delta_{js} + \delta_t + \varepsilon_{jst}, \quad (3.2)$$

where all variables and specification details are defined as in Equation (3.1).

Results from Equation (3.2) are shown in Figure 3.2(b). Coefficients in blue circles and yellow diamonds denote β_k^l and β_k^h estimates, respectively. Coefficients in light gray squares denote β_k coefficients from Equation (3.1). Products that consumers believed to be high-calorie (yellow diamonds) saw an initial drop in demand that faded 6 months after the policy implementation. In contrast, products consumers thought were relatively healthy but actually received a label (blue circles) saw a persistent decrease in demand of around 40%.[5] These empirical findings suggest that labels are especially effective for products about which consumers are more misinformed.

[5]The difference between the average value of $\hat{\beta}_t^l$ and $\hat{\beta}_t^h$ in the post-policy period is significant at the 98% confidence level.

Changes in nutritional content and prices

To study whether firms responded to the labeling policy by reformulating products, we compare the distribution of nutritional content before and after the policy was implemented. In 2016, 55 cereal products were above the threshold for caloric concentration. In 2018, 13 of those products reduced their concentration of calories to below the threshold, with eight of them bunching at the threshold of 350 kcal per 100 grams. We observe a similar pattern when we look at sugar concentration. In 2016, 27 regulated products were above the threshold. In 2018, 9 of these reduced their sugar content to be below the threshold and 6 reduced it to between 20 and 22.5 grams of sugar per 100 grams of cereal (see Appendix C.4, Figure C.4.1). This suggests that firms chose to respond strategically to the labeling policy, bunching at the threshold to avoid receiving a label.

This bunching results in a net reduction in the caloric and sugar concentration of cereal products offered in the market. The weighted average of the caloric concentration of products decreased from 383.6 to 372.8 kcal per 100 grams, while the weighted average of the sugar concentration of products decreased from 21.54 to 19.06 grams of sugar per 100 grams of cereal; weights are assigned by pre-policy revenue.

In Appendix C.1, we show that labeled products saw an average decrease of 5.5% in prices relative to unlabeled products. This may be explained by a combination of firms increasing markups on unlabeled products that now face higher demand (and vice versa) and an increase in marginal costs of unlabeled products due to reformulation. We find no evidence of firms responding by changing product assortment or package size.

3.4 Demand for Breakfast Cereal

We now develop and estimate a model of supply and demand for cereal that can explain the descriptive facts presented above. We use the model to answer policy-relevant questions such as what the total effect of the policy was in terms of consumer welfare and per capita nutritional intake, where the optimal threshold should be set, and how warning labels compare with sugar taxes.

Demand model

Our demand model consists of a continuum of risk-neutral consumers, indexed by $i \in \mathcal{I}$, who are divided into two bins defined by being above or below the median household income in our sample. We refer to them as low- and high-SES consumers and denote them by their type $b \in \{l, h\}$. We refer to each store-period combination as a "market" and index it by t. There are J products indexed by $j \in \mathcal{J}$ and one outside good (i.e. the option to buy no product). Each product j is produced by a firm $f \in \mathcal{F}$ and characterized by (r_j, p_{jt}, w_{jt}), where r_j is is a vector of indicator variables denoting the subcategory the product belongs to (plain, sugary, chocolate, granola, oatmeal); p_{jt} is its price in market t; and w_{jt} is its vector of nutritional content.

Our model departs from the standard random coefficients demand model (e.g., Berry, Levinsohn, and Pakes, 1995; Nevo, 2001) in an important way. We allow the nutritional content, w_{jt}, to affect utility through the negative long-run health consequences of consuming unhealthy goods. Nevertheless, because nutritional content may not be directly observed by consumers, their choices are based on their beliefs about it. As a consequence, consumer choices do not necessarily maximize consumer utility, which leaves space for government interventions with the potential to improve consumer welfare.

We assume that the utility derived by individual i when purchasing product j can be split into three main components:

$$u_{ijt} = \underbrace{\delta_{ijt}}_{\text{experience/taste}} - \underbrace{\alpha_i p_{jt}}_{\text{price paid}} - \underbrace{w'_{jt}\phi_i}_{\text{health consequences}} . \tag{3.3}$$

The first component, denoted by δ_{ijt}, corresponds to the aspect of utility that comes from the experience of consuming product j and is assumed to be observed by consumers when making the decision to buy the product. It is a function of the product's characteristics (e.g., sweetness, mouthfeel, smell) and other individual-level and time-varying demand shocks (e.g., idiosyncratic preferences for some products, hunger relief, food craving). In particular, we assume that

$$\delta_{ijt} = r'_j\beta_i + \delta_{jb} + \delta_{T(t)b} + \delta_{S(t)b} + \xi_{jtb} + \epsilon_{ijt}, \tag{3.4}$$

where β_i represents individual preferences for different subcategories; δ_{jb}, $\delta_{T(t)b}$, and $\delta_{S(t)b}$ are product, period, and store fixed effects, respectively, all specific to each consumer type; and ξ_{jtb} is a product-market-type specific idiosyncratic demand shock. ϵ_{ijt} is a consumer-specific demand shock that jointly follows a generalized extreme value distribution that follows the distributional assumptions of a one-nest nested logit model, where all inside goods are in the same nest. We denote the intra-nest correlation by ρ. We assume that $\beta_i \sim \mathcal{N}(0, \Sigma_\beta)$.

Note that this model specification does not allow the experience aspect of the utility to vary with changes in nutritional content, w_{jt}. As we will discuss later, we restrict firms to reformulations that maintain the taste of products constant. In other words, when changing w_{jt}, firms replace critical nutrients with alternative ingredients that maintain the sweetness, mouthfeel, smell, and other perceivable attributes.

The second element in the utility function, $\alpha_i p_{jt}$, corresponds to the disutility derived from paying price p_{jt} for product j. The parameter $\alpha_i \sim \log\mathcal{N}(\alpha_b, \sigma_\alpha)$ governs the price elasticity.

Finally, $w'_{jt}\phi_i$ corresponds to the negative long-term health consequences of consuming unhealthy products. The parameter $\phi_i \sim \log\mathcal{N}(\phi_b, \Sigma_\phi)$ represents the marginal damage perceived by consumer i from consuming additional critical nutrients w_{jt}.[6] Consumers do not know the true nutritional content, w_{jt}, but have prior beliefs, π_{ij}, about it. We assume

[6] Note that ϕ_i does not need to be the same for consumers and the social planner. So far, we are mostly interested in modeling consumer behavior. In Section 3.6, in which we discuss the normative implications of

that prior beliefs, π_{ij}, follow a normal distribution $\mathcal{N}(\mu_{jb}, \Omega_{jb})$. This allows both moments of the beliefs distribution to vary across products and consumer type. Additionally, we assume that the non-diagonal elements of Ω_{jb} are zero. This implies that sugar labels do not change beliefs about calories and vice versa.

Based on their beliefs, consumer i chooses the product that maximizes their expected utility:

$$\mathbb{E}_{\pi_{ij}}[u_{ijt}] = \delta_{ijt} - \alpha_i p_{jt} - \mathbb{E}_{\pi_{ij}}[w_{jt}|L_{jt}]'\phi_i, \tag{3.5}$$

where $\mathbb{E}_{\pi_{ij}}$ denotes the expectation operator over prior beliefs π_{ij} and $L_{jt} \in \{\text{pre-policy}, \text{no}, \text{yes}\}$ denotes the label status of product j in market t. We assume that consumers form their beliefs by using the observed labels (or lack thereof) and applying Bayes' rule.[7]

We denote the set of consumers that choose product j in market t by

$$\Theta_{jt} = \{i \in \mathcal{I}_t : \mathbb{E}_{\pi_{ij}}[u_{ijt}] \geq \mathbb{E}_{\pi_{ki}}[u_{ikt}], \ \ \forall k \in \mathcal{J}_t\}, \tag{3.6}$$

where $\mathcal{J}_t$ is the set of products available in market t, which includes the outside good, and $\mathcal{I}_t$ is the set of consumers who shop at least one time in supermarket $S(t)$, which we normalize to have mass one. The market share of product j in market t is given by $s_{jt} = \int_{i \in \Theta_{jt}} di$, while the share of consumers of type b who prefer product j in market t is given by $s_{jtb} = \int_{i \in \Theta_{jt} \cap b} di / \int_{i \in b} di$.

Modeling beliefs in our setting is essential. A model that ignores beliefs and in which labels enter into the utility function directly can lead to misleading conclusions. Only including a label dummy for the post-policy period would not capture the heterogeneity in responses that we observe in Figure 3.2(b). In the cereal market, products with a high-in-sugar label are also products that were already known to be high in calories and sugar. As a result, the products most affected by the policy were those that got a high-in-calories label but not a high-in-sugar one and were believed to be low in calories. A model that assumes beliefs away would have interpreted this result as consumers disliking high-in-calorie labels but liking high-in-sugar ones. Once we consider beliefs, we find that consumers dislike high concentrations of both calories and sugar. Not fully capturing the effects in demand would also lead to misleading incentives from the supply side when choosing which products to reformulate.

In Appendix C.2, we explore the implications of the main assumptions embedded in our demand model. We investigate the importance of using a static model, excluding salience

the model, we extend it to accommodate additional market imperfections such as lack of self-control or time inconsistency.

[7] We assume that consumers do not take into account product reformulation. We make this assumption for two reasons. First, interviews with consumers in Chile suggest that they did not realize that products may be bunching at the regulatory nutritional thresholds. Second, this assumption simplifies the calculation of consumers' posteriors and the solution of the market equilibrium.

effects, assuming invariant taste, and disregarding advertisement effects. We justify these modeling decisions and show that our primary findings are robust to modifying these assumptions.

Estimation and identification

To estimate the model, we aggregate the data at the product-store-period-consumer-type level. We estimate the model using the generalized method of moments proposed by Berry et al. (1995), but fixing consumer-type-level shares, s_{jtb}, at the observed levels. The estimating moment conditions are given by $\mathbb{E}[\xi_{jtb} Z_{jtb}] = 0$, where ξ_{jtb} is the demand shock from Equation (3.4) and Z_{jtb} are instruments that we describe below. We now discuss what variation in the data identifies each parameter and what instruments we use to exploit such variation.

Price coefficient

To identify α_b, the first moment of the price coefficient, we construct simulated instruments using the price of cereal inputs (Backus, Conlon, and Sinkinson, 2021). We collected the ingredients list of each cereal product, with the corresponding percentages of the main ingredients on them (e.g., Cheerios has 29% of corn, 21% of wheat, and 8% of oats), and combined it with historical price data on commodities from www.nasdaq.com to run the following regression:

$$p_{jt} = \sum_k \beta_k \upsilon_{kt} \varsigma_{kj} + d_j + d_{T(t)} + d_{S(t)} + \eta_{jt}, \tag{3.7}$$

where υ_{kt} is the price of commodity k in period $T(t)$ and ς_{kj} is the share of commodity k contained in product j in the pre-policy period. We include product, period, and store fixed effects. Commodities are corn, wheat, and oats. We then construct a price predictor given by

$$\hat{p}_{jt} = \sum_k \hat{\beta}_k \upsilon_{kt} \varsigma_{kj} + \hat{d}_j + \hat{d}_{T(t)} + \hat{d}_{S(t)}. \tag{3.8}$$

We use $\hat{p}_{jt}$ as an instrument for p_{jt}. It captures changes in prices that come from changes in commodity prices, and that are orthogonal to unobserved changes in demand. Since α_b takes different values for each consumer type, we interact the instrument with a consumer-type dummy.

Preferences for beliefs about health consequences

The identification of ϕ_i, the preferences over the perceived health consequences of consuming sugar and calories, and (μ_{jb}, Ω_{jb}), the parameters that govern the distribution of beliefs, is more difficult. In order to separate beliefs from preferences, we use information from the survey. We assume that the responses collected by the beliefs survey are informative about the ranking of and relative distance between μ_{jb} and μ_{kb}—the first moment of beliefs about the nutritional content of two different products—but that their absolute levels may be wrong.[8] We allow for the first moment of beliefs to be determined by $\mu_{jb} = \tilde{\mu}_{jb} + \mu$, where $\tilde{\mu}_{jb}$ is the average survey response regarding the expected value of nutritional content of product j among consumers of type b, and μ is a free parameter in our model that shifts the expected value of the nutritional content of all products among all consumers by a constant amount.[9]

We take Ω_{jb}, the second moment of beliefs about the nutritional content of each product, directly from the answers on the survey.

Combining the responses from the survey with the Bayesian model adds enough structure to jointly identify ϕ_b and μ. Figures 3.2 and 3.3 provide the intuition behind our identification strategy. To explain it, we illustrate the model prediction of changes in expected utility for two products, h and k (with $\tilde{\mu}_{hb} > \tilde{\mu}_{kb}$), at two parameter values, $\mu = \mu_1$ and $\mu = \mu_2$ (with $\mu_1 > \mu_2$).

In Figure 3.2, we plot the distribution of prior and posterior beliefs for products h and k conditional on not receiving a label. For ease of exposition, we assume that $\Omega_h = \Omega_k$. In panels (a) and (b), we plot beliefs when $\mu = \mu_1$, and in panels (c) and (d) when $\mu = \mu_2$. To recover posterior beliefs (dashed lines), we truncate prior beliefs at the policy threshold, which is invariant to μ. We denote the absolute change in the expected value of w_j induced by the labeling policy at parameter value μ by $\Delta\,\mathbb{E}^{\mu}[w_j|L_j]$, where $j = \{h,k\}$. Intuitively, $\Delta\,\mathbb{E}^{\mu_1}[w_j|L_j] > \Delta\,\mathbb{E}^{\mu_2}[w_j|L_j]$ for $j = \{h,k\}$ when $\mu_1 > \mu_2$. Moreover, $\Delta\,\mathbb{E}^{\mu_1}[w_h|L_h] - \Delta\,\mathbb{E}^{\mu_2}[w_h|L_h] > \Delta\,\mathbb{E}^{\mu_1}[w_k|L_k] - \Delta\,\mathbb{E}^{\mu_2}[w_k|L_k]$ for all (h,k) such that $\tilde{\mu}_{hb} > \tilde{\mu}_{kb}$. This nonlinear behavior of $\Delta\,\mathbb{E}^{\mu}[w_j|L_j]$ with respect to $\tilde{\mu}_{jb}$ and μ allows us to identify μ separately from ϕ_b.

We use Figure 3.3 to illustrate how the nonlinearity of $\Delta\,\mathbb{E}^{\mu}[w_j|L_j]$ with respect to $\tilde{\mu}_{jb}$ and μ helps us identify these parameters. The figure shows the change in expected utility from consuming product j as a function of $\tilde{\mu}_{jb}$. The solid line corresponds to $\mu = \mu_1$ and the dashed line to $\mu = \mu_2$. Different values of μ have different implications for the relative difference between the change in expected utility of products h and k. For large values of μ, the increase in expected utility from consuming product h will be larger than that from consuming product k. For small values of μ, the increase in expected utility will be small and similar for the two products.

Changes in expected utility present a kink-like structure, where μ determines the position

[8] We rely on the survey data for information on the relative levels, but not on the absolute levels of believed nutritional content of each product. We piloted three different survey designs, varying the reference products shown to respondents. We found that the levels of consumer responses were sensitive to the choice of the reference points, but the ranking and relative distance between answers for different products were robust across the survey designs.

[9] We normalize the elements of $\tilde{\mu}_b$ to have mean zero and the same variance as w^{pre} across products. The normalization implies that, in terms of changes in expected utility, a change in beliefs of 1 standard deviation is equivalent to a change in nutritional content of 1 standard deviation if nutritional content was observed. μ is measured in standard deviations and is constant for both nutrients.

of the kink in the $\tilde{\mu}_{jb}$ space. All unlabeled products to the left of the kink will experience small changes in expected utility. All unlabeled products to the right of the kink will experience an increase in expected utility. For products to the right of the kink, the increase in expected utility will be larger when $\tilde{\mu}_{jb}$ is higher. The differential change in expected utility between products implies a differential change in observed market shares. The shape of the change in observed market shares will identify the position of the kink and, therefore, the value of μ. The parameter ϕ_b, on the other hand, will determine the rate at which the change in expected utility increases with $\tilde{\mu}_{jb}$, which is given by the slope of the right side of the curve in Figure 3.3. Thus, ϕ_b will be identified by the relative differences in the changes of observed demand between products on the right side of the kink.

To bring this to the data, we first construct a predictor, $\hat{L}_{jt}$, of whether a product gets labeled or not that is uncorrelated with potential demand shocks, ξ_{jtb}. The predictor uses the cereal categories r_j and the pre-policy nutritional content as inputs, and estimates a random forest model to avoid overfitting. Distance from the policy threshold in the pre-policy period and heterogeneity in the cost of departing from the threshold driven by r_j explain most of the bunching, which provides us with an instrument that is highly correlated with labeling status. We then split products into different bins based on answers on the survey regarding the first moments of beliefs, $\tilde{\mu}_{jb}$. We denote these bins by B_μ. As illustrated in Figure 3.3, the model provides sharp predictions about how demand should change as a function of prior beliefs μ_{jb} and label status L_{jt}. By minimizing the moments $\mathbb{E}[\hat{L}_{jt} \times B_\mu \times \hat{\xi}_{jtb}]$, we impose conditions over $\hat{\xi}_{jtb}$ that prevent the patterns in Figure 3.3 from being explained by differential demand shocks. Without these moment restrictions, our model could explain the fact that products believed to be low in calories but which received a high-in-calories label experienced a reduction in demand, by assigning negative demand shocks to such products in the post-policy period. These moment conditions prevent such distribution of shocks, and thus identify ϕ_b and μ.

Preference heterogeneity

Finally, we need to identify Σ_β, σ_α, Σ_ϕ, and ρ, which are the parameters that govern the substitution patterns between different products and to the outside good. To do so, we construct three sets of market-level instruments. The first two sets of instruments exploit changes in competitors' cost-shifters, which through changes in prices should shift the probability that consumers substitute from one product to the other. The third set of instruments exploits the entrance of new products to the market that induce changes in the competitive environment. Let τ_{jt} be the first time a given product enters supermarket $S(t)$. Then, the three set of instruments are given by

$$z_t^{r,1} = \underset{j \in r,t}{\text{mean}}\{\hat{p}_{jt}\}, \qquad z_t^{r,2} = \underset{j \in r,t}{\text{pctile}^{20,80}}\{\hat{p}_{jt}\}., \qquad z_t^{r,3} = \sum_{j \in r,t} \mathbb{1}\{t \geq \tau_{jt}\}.$$

The first set of instruments corresponds to the average price predictor of all products in each cereal category r and market t. The intuition behind the instrument is that when commodities usually used in a given subcategory, r, are cheap, consumers will be more likely to substitute toward products in that subcategory. For example, if oat prices in a given period are low, we should expect to see more substitution toward oat products in that period.

The second set of instruments corresponds to the 20th and 80th percentiles of the price predictor among all products in a given cereal category r and market t. These instruments work in a fashion similar to the first set of instruments, but add additional moments of the predicted price distribution of competitors' products, which increases statistical power.

The third set of instruments exploits the timing of the entrance of different brands into different stores. These instruments measure the total number of products from each subcategory, r, that have ever entered store $S(t)$ before period $T(t)$. The identifying assumption is that the first entry of a product at the supermarket level is not correlated with demand shocks. We believe this is a reasonable assumption given that Walmart is increasing its assortment in many product categories, including cereal (see Appendix C.1). At the beginning of the sample period, there are on average 52 products available in each market. By the end of the sample period, the average number of available products per market grows up to 73 products. Empirically, the increase in product assortment is not correlated with the timing of the policy. The intuition behind the instruments is that when more products are available and variety increases, consumers are less likely to substitute toward the outside option, which helps us to identify ρ.

Results

Our estimated demand parameters are presented in Table 3.1. Our estimates imply an average own-price elasticity of −3.1, with a higher absolute elasticity among low SES households (−3.33 vs. −2.74). We also find that products in the same subcategory, r_j, are closer substitutes. We present the matrix of own- and cross-price elasticities of the most important products from each subcategory in Appendix C.4, Table C.4.1. These elasticities imply median markups—defined as the ratio of price minus marginal cost to price—of 46% in the pre-policy period.[10] These results are similar to those in previous papers that estimate demand for cereal in the U.S. market and find elasticities between −2.3 and −4.3 and median markups of 34%-42% (Nevo, 2001; Michel and Weiergraeber, 2018; Backus et al., 2021). Our estimates are also comparable to accounting estimates provided by the Chilean antitrust agency, which estimates markups of 45% for the largest cereal brand in Chile (FNE, 2014).

The estimates for ϕ_i indicate that an average consumer is willing to pay 9.9% and 7.6% of the average price of cereal to reduce the sugar and caloric concentration of products, respectively, by 1 standard deviation (12 grams of sugar and 25 kilocalories per 100 grams of cereal, respectively), while keeping the taste constant. For example, *Original Cheerios* contains 5 grams of sugar per 100 gram, while *Honey Nut Cheerios* contains 32.5 grams of sugar per 100 grams. According to our model, consumers would be willing to pay \$0.7 more

[10]We present the full distribution of markups in Appendix C.4, Figure C.4.2.

for a 550 grams family size box of *Honey Nut Cheerios* if it contained the sugar content of *Original Cheerios* but kept its own taste. In Figure 3.4, we show the distribution of willingness to pay among low- and high-SES consumers to reduce the sugar and caloric concentration of products by 1 standard deviation, while keeping the taste constant. We find substantial consumer heterogeneity, especially for preferences over sugar content.

We find an intra-nest correlation of $\rho = 0.96$, which suggests that there is little substitution from inside goods to the outside good. This should be taken with caution as it is larger than that estimated in the previous literature. However, we show in Appendix C.4, Figure C.4.3, that our main results are qualitatively similar when we impose a lower value of ρ.[11] Finally, μ shifts beliefs about sugar and caloric concentration by 0.13 standard deviations downward.[12]

3.5 Supply: Pricing and Nutritional Content

Supply model

Each firm f has a bundle of products $\mathcal{J}_f$ that it can produce. To produce a given product j, firms use two types of inputs: critical nutrients w_{jt} (e.g., sugar), and other inputs m_{jt} (e.g., sucralose, polyols).[13] The taste of a product depends on the concentration of these inputs and is given by a product-specific production function $\delta_j(w_{jt}, m_{jt})$. We restrict firms to reformulations that maintain the product's taste, $\bar{\delta}_j$, constant. That is, when firms reformulate their products, they choose inputs to always achieve the same level of sweetness, crunchiness, smell, etc. This is consistent with industry participants' descriptions of how reformulation was accomplished.[14] Since taste, $\bar{\delta}_j$, is invariant, firms need to choose w_{jt} and m_{jt} such that

$$\delta_j(w_{jt}, m_{jt}) = \bar{\delta}_j \tag{3.9}$$

The cost of producing a product depends on the nutritional content w_{jt}, other inputs m_{jt} and an additive cost-shifter ϑ_{jt}:

$$\tilde{c}_{jt}(w_{jt}, m_{jt}) \quad = \quad p_w w_{jt} + p_m m_{jt} + \vartheta_{jt}. \tag{3.10}$$

[11] At face value, the estimated substitution to the outside option would have unrealistic implications for how a monopolist in this market would behave. It could also affect the interpretation of our tax counterfactual as the overall demand for cereal would be insensitive to higher taxes.

[12] We plot the estimated values of μ_{jb} in Appendix C.4, Figure C.4.4. Regarding Ω_j, its diagonal elements range from 20-40$\left(\frac{\text{g}}{100\text{ g}}\right)^2$ for sugar and 200-325$\left(\frac{\text{kcal}}{100\text{ g}}\right)^2$ for calories.

[13] Note that other inputs, m_{jt}, might also have adverse health consequences. In our model, we let the policy-maker decide what nutrients are considered harmful (i.e., what nutrients are included in the vector w_{jt}) and assume all other inputs to be harmless.

[14] We interviewed the consumer product managers of the two largest cereal companies. They confirmed that an explicit goal of the reformulation process is that the new version of the product is indistinguishable from the previous one. To achieve this, firms follow several steps that include conducting expert focus groups and randomized blind tests.

From Equations (3.9) and (3.10) we can redefine the marginal cost of producing product j as

$$c_{jt}(w_{jt}) = p_w w_{jt} + p_m m_j(w_{jt}, \bar{\delta}_j) + \vartheta_{jt}, \tag{3.11}$$

where $m_j(w_{jt}, \bar{\delta}_j)$ is the inverse function of $\delta_j(w_{jt}, m_{jt})$ in Equation (3.9), provided that $\delta_j(w_{jt}, m_{jt})$ is invertible.

Let ν_j, which we will call the *bliss point* of product j, be the value of w_{jt} that minimizes marginal cost (i.e., ν_j is such that $\nabla c_{jt}(\nu_j) = 0$). The bliss point is an attribute of the product and corresponds to the concentration of critical nutrients that product j should have to achieve taste $\bar{\delta}_j$ at minimum cost. In the cereal market, for example, we should expect *Honey Nut Cheerios* to have a higher bliss point for sugar than *Original Cheerios*, since the former is a sweetened version of the latter.

Departing from the bliss point is possible but costly. For example, after the food labeling policy was introduced, firms in the breakfast cereal market replaced sugar with artificial alternatives such as sucralose and polyols.[15] This reformulation results in a more expensive product, captured in our model by the functional form of $c_{jt}(w_{jt})$. For each product, we approximate the marginal cost function by a second-order Taylor polynomial around the bliss point, such that

$$c_{jt}(w) = \underbrace{\bar{c}_{jt}}_{\text{baseline cost}} + \underbrace{(w - \nu_j)' \Lambda_j (w - \nu_j)}_{\text{change in cost due to reformulation}}, \tag{3.12}$$

where $\Lambda_j = \begin{bmatrix} \lambda_j^s & 0 \\ 0 & \lambda_j^c \end{bmatrix}$ with $\lambda_j^n > 0$ for $n \in \{s, c\}$ and all products j. We assume that λ_j^n is drawn from a lognormal distribution with parameters $(\mu_\lambda^n, \sigma_\lambda^n)$, where $\mu_\lambda^c = \bar{\mu}_\lambda^c + \vartheta_\lambda^c \nu_j^s$ while $\mu_\lambda^s = \bar{\mu}_\lambda^s$. This allows for the cost to reformulate calories to depend on the baseline sugar concentration of the product. However, having zeros on the non-diagonal elements of Λ_j implies that the costs of marginally reducing sugar and caloric concentration are not correlated. These assumptions are consistent with the data, where we find low correlation between caloric and sugar content and between changes in these induced by reformulation, but we find that high-in-sugar products were less likely to reformulate calories.

The firm's profit maximization problem is given by

$$\max_{\{p_{jt}, w_{jt}\}_{j \in \mathcal{J}_{ft}}} \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - c_{jt}(w_{jt})) \cdot s_{jt}(\mathbf{p_t}, \mathbb{E}_\pi[\mathbf{w_t} | \mathbf{L_t}]), \tag{3.13}$$

where s_{jt} is the market share of product j in market t, which depends on the vector of all prices $\mathbf{p_t}$ and all individuals' expectations about the nutritional content of all products in

[15] We collected data on specific ingredients of 17 out of the 20 products that reformulated in our sample. We found that after the policy is implemented, 47% start using maltitol (a type of polyols), 29% sucralose, and 35% stevia.

the market, $\mathbb{E}_{\pi}[\mathbf{w_t}|\mathbf{L_t}]$. In the absence of any government intervention, the firm chooses

$$w^*_{jt} = \nu_j \tag{3.14}$$
$$p^*_{jt} = c_{jt}(w^*_{jt}) + \Delta^{-1}_{(j,\cdot)}\mathbf{s_t}, \tag{3.15}$$

where the (j,k) element of Δ is given by

$$\Delta_{(j,k)} = \begin{cases} \frac{-\partial s_k}{\partial p_j} & \text{if } k \in \mathcal{J}_{ft} \\ 0 & \text{otherwise,} \end{cases} \tag{3.16}$$

and $\Delta^{-1}_{(j,\cdot)}$ is the jth column of the inverse of Δ. Equation (3.14) states that firms will choose the nutritional content of product j to be equal to its bliss point.[16] Equation (3.15) implies price-cost markups given by $\Delta^{-1}_{(j,\cdot)}\mathbf{s_t}$, where $\Delta^{-1}_{(j,\cdot)}$ takes into account that by increasing price j, demand for other products produced by firm f might increase.

When the food labeling regulation is in place, the demand function $s_{jt}(\mathbf{p_t}, \mathbb{E}_{\mathcal{I}_t}[\mathbf{w_t}|\mathbf{L_t}])$ becomes discontinuous in w_{jt} at the threshold. Firms have incentives to reduce the nutritional content of products whose bliss points are to the right of, but close to, the threshold. By marginally increasing the production cost of a product close to the threshold, firms can choose w_{jt} to be right below the threshold, thus changing consumers' conditional expectations and inducing large increases in demand. This explains the bunching observed in the data.

In Appendix C.3, we explore the implications of the main assumptions embedded in our supply model. We study the importance of the firms choices' timing in choosing prices and nutritional content and of assuming that reformulation does not change the taste of products but increases marginal cost. We justify these modeling decisions and show that our primary findings are robust to modifying these assumptions.

Estimation

To estimate the supply model, we need to recover three key parameters: (a) the marginal cost of producing a product in the absence of reformulation, $\bar{c}_{jt}$, (b) the products' bliss points, ν_j, and (c) the cost of reformulating, Λ_j, which is determined by $\left(\bar{\mu}^n_{\lambda}, \sigma^n_{\lambda}, \vartheta_{\lambda}\right)$.

We recover $c_{jt}(w^*_{jt})$ and ν_j from the firm's first-order conditions (Equations (3.14) and (3.15)). We then estimate $\bar{\mu}^n_{\lambda}$, σ^n_{λ}, and ϑ_{λ} by exploiting variation in firms' decisions to bunch.

Using our demand estimates, we compute the equilibrium at the current parameters and labels. We then ask, for each product, what would be the value of λ^n_j that would render firm $f(j)$ indifferent between choosing the bliss point level ν^n_j or having product j bunching at the threshold, keeping all other products' nutritional content decisions fixed. We denote the

[16] In the absence of any policy, demand does not depend on w_{jt} or m_{jt}. In that case, the firm's optimal decision is to choose a combination of w_{jt} and m_{jt} that minimizes marginal cost.

indifference value by $\tilde{\lambda}_j^n$. Then, the probability that product j bunches in nutrient n is given by $P_{B_j^n} = Pr(\lambda_j^n \leq \tilde{\lambda}_j^n)$.[17]

We estimate $(\mu_\lambda^n, \sigma_\lambda^n)$ for $n \in \{s, c\}$ and ϑ_λ via GMM by imposing that the difference between the probability of bunching, $P_{B_j^n}$, and whether a product bunches or not, B_j^n, has mean zero and is uncorrelated with the product's bliss point ν_j:

$$\begin{aligned}
\mathbb{E}[(B_j^n - P_{B_j^n})] &= 0 \quad \text{for } n \in \{s, c\} \\
\mathbb{E}[(B_j^n - P_{B_j^n})\nu_j^n] &= 0 \quad \text{for } n \in \{s, c\} \\
\mathbb{E}[(B_j^c - P_{B_j^c})\nu_j^s] &= 0.
\end{aligned}$$

Once we estimate $(\mu_\lambda^n, \sigma_\lambda^n)$, we calculate $\bar{c}_{jt}$ by solving

$$\bar{c}_{jt} = c_{jt}(w_{jt}) - \mathbb{E}_\lambda[(w_{jt} - \nu_j)'\Lambda_j(w_{jt} - \nu_j)|B_j]. \tag{3.17}$$

Results

Our estimated supply parameters are presented in Table 3.2. To interpret these parameters, we calculate $\mathbb{E}[\lambda_j^n | B_j^n = 1]$, the expected value of λ_j^n conditional on product j bunching in nutrient n. We find an average value of $0.151 \frac{\text{¢}}{(\text{g}/100\text{ g})^2}$ in the case of sugar and $0.016 \frac{\text{¢}}{(\text{kcal}/100\text{ g})^2}$ in the case of calories. The average reduction in sugar concentration among products bunching in sugar is 8.2 grams per 100 grams, while the average reduction in caloric concentration among products bunching in calories is 24.9 kilocalories per 100 grams. Putting everything together, our model finds that the average expected increase in marginal cost for products bunching in any nutrient is 2.8¢ per 100 grams, which is equivalent to 4.4% of the average price of cereal.

To assess the accuracy of our estimates, we run a regression to calculate how our estimates of marginal cost, $c_{jt}(w_{jt}^*)$, differ between products that did and did not bunch at nutritional thresholds and compare them with the change in marginal cost implied by our estimated supply parameters that govern Λ_j. To do this, we estimate the following equation:

$$c_{jt}(w_{jt}^*) = \beta \cdot B_j \cdot Post_t + \delta_{js} + \delta_t + \varepsilon_{jt}, \tag{3.18}$$

where $c_{jt}(w_{jt}^*)$ is computed using the firm's first-order conditions, B_j is a dummy indicating whether product j is bunching in the post-period, and δ_{js} and δ_t are product-store and period fixed effects, respectively. The estimated coefficient $\hat{\beta}$ from Equation (3.18) suggests an average change in marginal cost of 3.1¢ per 100 grams, slightly larger than the 2.8¢ per 100 grams derived from Equation (3.17) of our model.

[17] Note that λ_j^n is not point-identified. From the data, we learn that for products bunching in nutrient n, $\lambda_j^n \leq \tilde{\lambda}_j^n$, and that for products not bunching in nutrient n, $\lambda_j^n > \tilde{\lambda}_j^n$. However, we cannot recover the exact value of λ_j^n. Treating λ_j^n as a random coefficient drawn from a known distribution allows us to overcome this identification problem.

We also compare the model-based predicted probability of each product bunching in a given nutrient with what actually happened in the data. Figure 3.5 shows the probability of bunching predicted by the model for each product to the right of the policy threshold. Products in gray are products that bunched in the data and did not receive a label. Products in color are those that did not bunch. The model predicts correctly that products ex ante closer to the threshold are more likely to bunch.

In Appendix C.4, Figure C.4.5, we also show that our model correctly predicts that products for which prior beliefs about nutritional content were lower have a higher probability to bunch.

3.6 The Impact of Food Labeling Policies

In this section, we use our model to evaluate the effects of food labeling policies on nutritional intake and overall welfare. We start by simulating the Chilean Food Act under several counterfactuals that isolate different economic forces. We then study optimal policy design and compare food labels with sugar taxes, which is the most prominent alternative policy instrument.

Equilibrium effects of food labels

We estimate the effects of the Chilean Food Act on consumer choices, firms' production and pricing decisions, nutritional intake, and consumer welfare. To disentangle the roles of demand and supply in changes in nutritional intake and consumer welfare, we run four counterfactuals. The first counterfactual, denoted by (0), *no intervention*, corresponds to the case in which no policy is in place. To isolate demand forces, we compare the no-intervention benchmark with a situation in which products receive labels according to the regulatory thresholds and suppliers are not allowed to respond. We denote this counterfactual by (1), *demand only*. We then compute counterfactual (2), *price response*, in which—in addition to receiving labels—we allow suppliers to optimally choose prices while keeping nutritional content constant. We use counterfactual (2) to measure additional changes in consumer welfare driven by competition and product differentiation, which can either decrease or increase prices. The differences in consumer welfare between (1) and (2) are thus ambiguous. Finally, we compute counterfactual (3), *equilibrium*, in which we also allow firms to change the nutritional content of their products. This corresponds to the equilibrium model presented in Sections 3.4 and 3.5. The expected change in consumer welfare from counterfactual (2) to (3) is also ambiguous. Although firms improve product quality by reducing the concentration of critical nutrients, production costs increase, which leads to higher prices for consumers. Whether the policy under counterfactual (3) increases or decreases consumer welfare relative to (0) is therefore an empirical question.

To estimate consumer welfare, we cannot use a standard revealed preferences approach, because in our setting consumer choices do not necessarily maximize utility. We follow Allcott (2013), who offers a framework to calculate consumer welfare in situations in which consumers' ex ante expected utility differs from what they actually experience when consuming their chosen alternative. To do so, we define consumers' utility from the perspective of the social planner as

$$u_{ijt}^{SP} = \delta_{ijt} - \alpha_i p_{jt} - w'_{jt}\phi_i\lambda. \tag{3.19}$$

The social planner's utility from Equation (3.19) differs from the expected utility function consumers use to make choices in Equation (3.5) in two different ways. First, the social planner's utility depends on the true nutritional intake w_{jt} rather than the expected one. Second, we allow the social planner to disagree with consumers about the marginal damage of consuming additional critical nutrients by multiplying ϕ_i by a constant λ. This allows our model to accommodate additional market imperfections, such as externalities in the form of financial health-care costs or internalities in the form of self-control problems, time-inconsistency, or misperceptions about the individual damage caused by critical nutrients, ϕ_i. For the main part of our analysis, unless otherwise stated, we focus on results for the case in which $\lambda = 1$ (i.e., in which there are no additional market imperfections). Equation (3.19) makes specific normative assumptions and does not allow, for example, for models in which "ignorance is bliss" (i.e., consumers are better off not knowing that they are engaging in harmful behavior) or in which labels affect utility in some other way.[18]

Average consumer welfare in market t under counterfactual (x) is given by

$$CW^t(x) = \sum_j \left\{ \int_{\Theta_{jt}^{(x)}} \frac{1}{\alpha_i}(\delta_{ijt} - \alpha_i p_{jt}^{(x)} - w_{jt}^{(x)}\phi_i\lambda)di \right\},$$

where $p_{jt}^{(x)}$ and $w_{jt}^{(x)}$ are the price and nutritional content of product j in market t in counterfactual (x). $\Theta_{jt}^{(x)}$ is the set of consumers who prefer product j in counterfactual (x). Since taste is constant, δ_{ijt} does not vary across counterfactuals. The total mass of potential consumers is normalized to be one in each market. We present the average change in consumer welfare between counterfactuals (x) and (0) in Figure 3.6, and decompose it between how much of it is driven by changes in nutritional intake, changes in dollars spent, and changes in the average taste of products that are consumed.

We find that moving from a counterfactual with no intervention, (0), to one in which products get labeled but suppliers do not respond, (1), increases average consumer welfare by \$0.27 a year. This corresponds to 1.1% of the average yearly expenditure on cereal products. In the absence of supply-side responses, consumers shift demand from products high in critical nutrients to those low in critical nutrients. Since in the breakfast cereal market caloric and sugar content are positively correlated with prices, consumers end up

[18]Readers who disagree with this normative model can take home the positive results of our model: the changes in nutritional intake, the changes in dollars spent by consumers, and the changes in the taste of the products consumers choose. The normative model just adds weights to these positive results to aggregate them into a single index we call welfare.

consuming products that are cheaper but, according to the model, have lower taste (e.g., oatmeal).

We then allow firms to optimally set prices in response to the policy by simulating counterfactual (2). Under this counterfactual, we find that prices of unlabeled products go up while prices of labeled products go down. Overall, prices increase by 0.05% on average and gains in consumer welfare relative to counterfactual (0) are \$0.25 a year per capita (7% lower than under counterfactual (1)).

Under counterfactual (3), firms not only choose prices, but also the nutritional content of their products. We find large gains in consumer welfare from reducing caloric intake, mostly driven by products that become healthier due to reformulation.[19] Gains in consumer welfare due to lower intake of critical nutrients are 30% larger than under counterfactual (1). However, reformulation increases production costs, which leads to higher prices. The net effect is an average gain in consumer welfare of \$0.46 a year under counterfactual (3), which is 70% larger than under counterfactual (1).

On the firm side, average yearly profits per capita increase by only \$0.01, with substantial heterogeneity across firms. While some firms increased their profits by around 10%, others lost more than 20%. Who wins and who loses is closely related to how labels shift consumer beliefs. Firms with products that were believed to be healthy but ended up labeled experience the highest losses. This may explain why some firms opposed the Chilean Food Act so strongly when it was first implemented.

Finally, we consider an additional counterfactual in which consumers are perfectly informed about the nutritional content of products. This exercise informs us about the total welfare losses due to lack of information in the cereal market, and allows us to assess how well food labels approximate the best-case scenario of perfect information. We find that the food labeling policy achieves 8% of the consumer welfare gains that would be obtained under the perfect information counterfactual.

The design of food labeling policies

We now study the design of food labeling policies. We take the binary-signal structure of the policy as given, and study how nutritional intake and consumer welfare vary under different regulatory thresholds. Intuitively, in the absence of supply-side effects, thresholds should be set such that labels' informativeness is maximized. When supply-side responses are considered, policymakers can choose a different regulatory threshold that induces larger reductions in critical nutrients. To clarify the analysis, we simplify our model to only allow misinformation regarding sugar content.[20]

[19]Changes in consumer welfare from reducing sugar intake are negative. On one hand, firms reformulate products to have a lower concentration of sugar. On the other hand, more products are unlabeled in counterfactual (3), which means that the average sugar concentration among unlabeled products is higher. The latter effect offsets the potential benefits of the former effect.

[20]We assume consumers are perfectly informed about the nutritional content of calories in all counterfactuals.

We focus our analysis on counterfactuals (1), demand-only responses, and (3), the equilibrium model. Figure 3.8(a) shows the gains in consumer welfare under counterfactuals (1) and (3) for different policy thresholds. A naive policymaker who seeks to maximize consumer welfare but ignores equilibrium effects would set the policy threshold at 16.5 grams per 100 grams, the value at which consumer welfare is maximized under counterfactual (1). Consumer welfare under counterfactual (3), however, is maximized at 8.5 grams per 100 grams, at which point it is 20% larger than under the naive threshold.

Food labels vs. sugar taxes

We exploit the richness of our model to compare the effectiveness of food labels against sin taxes. We focus on sugar taxes, a widespread policy used in more than 40 countries (Allcott, Lockwood, and Taubinsky, 2019b). Most sugar taxes are structured as a per-ounce tax on any product with added sugar. However, Allcott et al. (2019b) recommend using tax designs that depend on the amount of sugar instead of the amount of product, to encourage consumers to switch to lower-sugar products and producers to reduce sugar content. We follow this tax structure. We assume that consumers observe the final after-tax price of products and cannot infer the concentration of critical nutrients by looking at prices. This is a reasonable assumption in our context, since sales taxes are not observed by consumers in Chile. We use ψ to denote the marginal value of public funds. To calculate consumer welfare, we distribute the tax money to consumers through a lump sum transfer (i.e., $\psi = 1$).

Extending the model from Section 3.5 to include sugar taxes, the firm's problem is given by

$$\max_{\{p_{jt}, w_{jt}\}_{j \in \mathfrak{I}_j}} \sum_{j \in \mathfrak{I}_j} (p_{jt} - c_{jt}(w_{jt}) - w_{jt}\tau) \cdot s_{jt}(\mathbf{p_t}, \mathrm{E}[\mathbf{w_t}])$$

where τ is the tax per gram of sugar and p_{jt} is the final price paid by consumers. From the first-order conditions, we have

$$\begin{aligned} \nabla c_{jt}(w^*_{jt}) &= -\tau \\ p^*_{jt} &= c_{jt}(w^*_{jt}) + \tau w^*_{jt} + \Delta^{-1}_{(j,\cdot)} \mathbf{s_t}, \end{aligned}$$

where the (j,k) element of Δ is given by equation (3.16). In this setting, firms have incentives to deviate from the bliss point, ν_j, and reduce the nutritional content of their products to pay lower taxes. Moreover, the price equation has an additional term given by the tax, which is proportional to the sugar content, and gets passed on to consumers through higher prices.

In Figure 3.8(b), we present gains in consumer welfare at different tax values. The optimal sugar tax (i.e., the tax that maximizes consumer welfare) is set at 0.3¢ per gram of sugar. This is not far from the value of sugar taxes implemented in some U.S. cities.[21] Gains in

[21] Philadelphia and Berkeley are the first two cities to pass a sugar tax in the U.S. In Berkeley, there is a 1¢ tax per ounce of sugar-sweetened beverages, equivalent to 0.32¢ per gram of sugar in the case of Coca-Cola.

consumer welfare with optimal sugar taxes are 29.5% lower than under food labels at the optimal policy threshold.

We find that taxes are 31% more effective at reducing sugar intake than food labels. However, they do this at a greater direct financial cost to consumers. Under the optimal tax level, consumers spend 2.6 additional dollars a year in taxes, equivalent to 7.5% of the total expenditure on cereal. Because taxes collected are relatively high, our results are sensitive to the choice of ψ, the marginal value of public funds.

Note that in contrast to food labels, sugar taxes are granular instruments, which are levied more heavily on products with higher levels of sugar. This is important for two reasons. First, sugar taxes have the potential to incentivize firms to reformulate all of their products in order to pay lower taxes, especially those with higher sugar content. Second, the effects of sugar taxes do not depend on consumers' beliefs. This makes taxes particularly appealing when λ, the parameter that accounts for additional market imperfections, is high.

Sensitivity to different values of λ and ψ

We take our values for λ from Allcott et al. (2019c), who estimate externalities from consuming sugar-sweetened beverages to be 0.8¢ per ounce, and internalities—which include the type of misinformation analyzed in this paper—to be around 1¢ per ounce. Taking into account that the median sugar-sweetened beverage has 3.25 grams of sugar per ounce, the additional marginal damage from consuming a gram of sugar is between 0.25¢ (only externalities) and 0.55¢ (externalities + internalities). In our model, this corresponds to $\lambda = 1.5$ and $\lambda = 2.1$, respectively.

The marginal value of public funds, ψ, can vary substantially depending on how tax money is spent. Hendren and Sprung-Keyser (2020) find that a large variety of policies targeted at adults in the United States have marginal values of public funds that range from $\psi = 0.8$ to $\psi = 1.2$.

In Figure 3.8(c), we show the values of λ and ψ for which labels are better than taxes and vice versa. Intuitively, larger values of λ favor taxes since they are better designed to deal with market imperfections not directly related to misinformation regarding w_{jt}. Taxes, however, impose a large burden on consumers who end up spending more on cereal. If the marginal value of public funds ψ is small, the resources collected through taxes will not contribute much to the total welfare. The smaller the value of ψ, the less effective taxes will be.

Heterogeneity in beliefs

In settings with heterogeneous agents, food labels can be more efficient than sugar taxes because their effects can be better targeted. To illustrate this point, consider a simple model in which half of the consumers have miscalibrated beliefs and the other half have accurate beliefs (i.e., $\mu_{jb} = \nu_j$, $\Omega_{jb} \to 0$). We call them uninformed and informed consumers,

In Philadelphia, the tax is 1.5¢ per ounce, equivalent to 0.48¢ per gram of sugar.

respectively. To gain intuition, let us focus on the case in which there are no supply-side responses. Ideally, the regulator would like to implement a targeted policy that only applies to uninformed consumers (e.g., food labels or sugar taxes for the uninformed population only). Although implementing a targeted policy is usually not possible, food labels will only affect the decisions of uninformed individuals and not those of consumers who are informed and were already making optimal choices, even when the instrument is not itself targeted. Taxes, on the other hand, are blunt instruments that generally change the actions of all consumers, and benefit some while hurting others.

Distributional consequences

The progressivity or regressivity of a policy depends on how the benefits (e.g., more information, correction of biases) and the costs (e.g., the burden of tax payments) vary across the income distribution. Two key parameters in our model are crucial in determining the incidence of each policy.

The first parameter is the extent to which low-SES consumers are more or less inclined than high-SES consumers to prefer products that are high in sugar. While food labels improve consumer welfare by providing information about the healthiness of products, taxes correct consumer behavior by inflating the prices of products that are high in sugar. If low-SES consumers prefer high-in-sugar products more than high-SES consumers do, then they will be charged disproportionately higher taxes. Depending on how the tax revenue is spent by the government, sugar taxes can benefit high-SES consumers relatively more. In the United States, for example, consumers with household incomes below $10,000 purchase 25% more grams of added sugar per calorie than do households with incomes above $100,000 (Allcott et al., 2019a). Sugar taxes are therefore more likely to be regressive than food labels.

The second parameter is the extent to which low-SES consumers are more or less informed than high-SES consumers regarding the nutritional content of products. An advantage of food labels relative to sugar taxes is that the former can be better targeted toward the uninformed population. Using survey data, Allcott et al. (2019c) find that U.S. consumers with household income below $10,000 score 0.82 standard deviations lower than consumers with household income above $100,000 on a nutrition knowledge questionnaire, which renders food labels more progressive than sugar taxes.

3.7 Conclusion

In this paper, we study the equilibrium effects of food labeling policies on nutritional intake and consumer welfare. Three key findings arise from our empirical analysis. First, the food labeling regulation caused consumers to substitute from labeled to unlabeled food products. Second, products that were perceived as healthy but received labels experienced the largest decline in demand. Third, suppliers responded to the policy by changing prices and reformulating their products.

We develop and estimate an equilibrium model of supply and demand for food and nutrients and use it to calculate the effects of food labeling policies on nutritional intake and consumer welfare. We find that food labels can be an effective way to improve diet quality and combat obesity. Our analysis shows that food labels are more effective when consumers have mistaken beliefs about products' healthiness, consumers value healthiness, reformulation that does not substantially change products' taste is feasible, and regulatory thresholds are set so that they provide useful information to consumers and encourage product reformulation.

We then use our model to compare food labels with sugar taxes. When compared with sugar taxes, food labels present both advantages and disadvantages. We show that food labels are more effective for tackling misinformation, but less effective for dealing with other market imperfections such as fiscal externalities, lack of self-control, or time inconsistency. Food labels are more progressive than sugar taxes, especially in settings in which the poor tend to consume more sugary products or in which the poor are more misinformed about the nutritional content of available products.

Our analysis shows how a theoretical framework combined with data can inform the design of policies to combat obesity by identifying and measuring the most relevant economic forces at work. Our model can accommodate a variety of settings and can be used to study the effects of food labels in categories other than cereal. It also provides a useful framework for comparing food labels with alternative policy instruments.

Food labels are a new and promising policy tool with the capacity to improve diet quality. While this paper covers important features of food labels, several unanswered questions remain. First, this paper focuses on a policy design in which labels act as a binary signal. New research suggests that more granular labels can be more effective in improving diet quality (Ravaioli, 2021). Second, food labels can incentivize firms to design new healthy products targeted to more informed consumers, which improves the bundle of available products in the long run. Finally, measuring long-run outcomes on health and wellbeing will be crucial in assessing the effectiveness of food labels.

Figures

Figure 3.1: Relative changes in equilibrium quantities

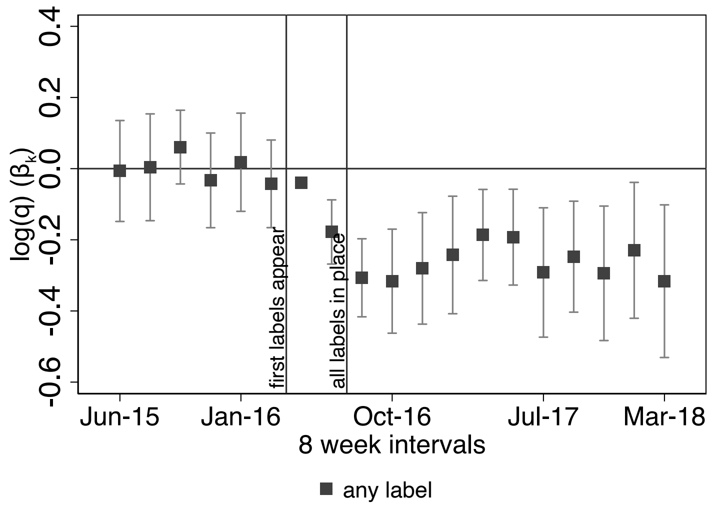

(a) By labeling status

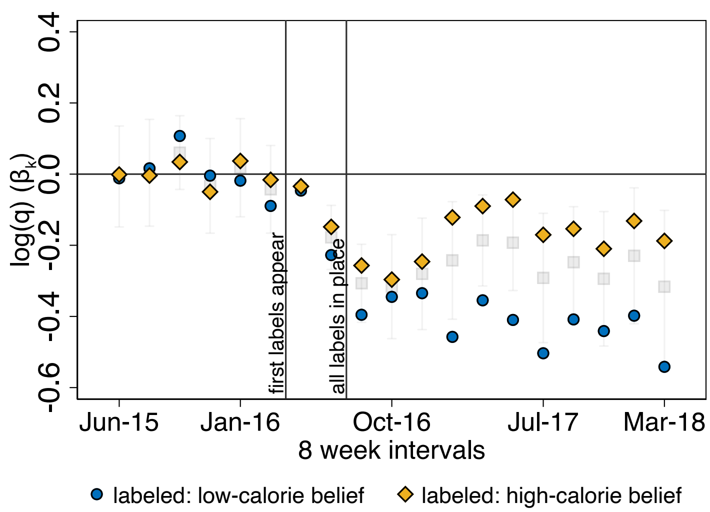

(b) By labeling status and prior beliefs

Notes: This figure presents the coefficients of our event study regressions. Panel (a) presents the β_k coefficients from Equation (3.1). Panel (b) displays the coefficients from Equation (3.2). Coefficients in blue circles, yellow diamonds, and light gray squares denote β_k^l, β_k^h, and β_k estimates, respectively. The vertical segments delimit the 95% confidence intervals. We run the regressions on the sample of 68 ready-to-eat cereals that show up in the pre- and post-policy periods. The sample consists of 27 unlabeled and 41 labeled products for a total of 194,510 observations.

Figure 3.2: Model-implied change in beliefs about about nutritional content, w, for products h and k at different values of μ upon not receiving a label

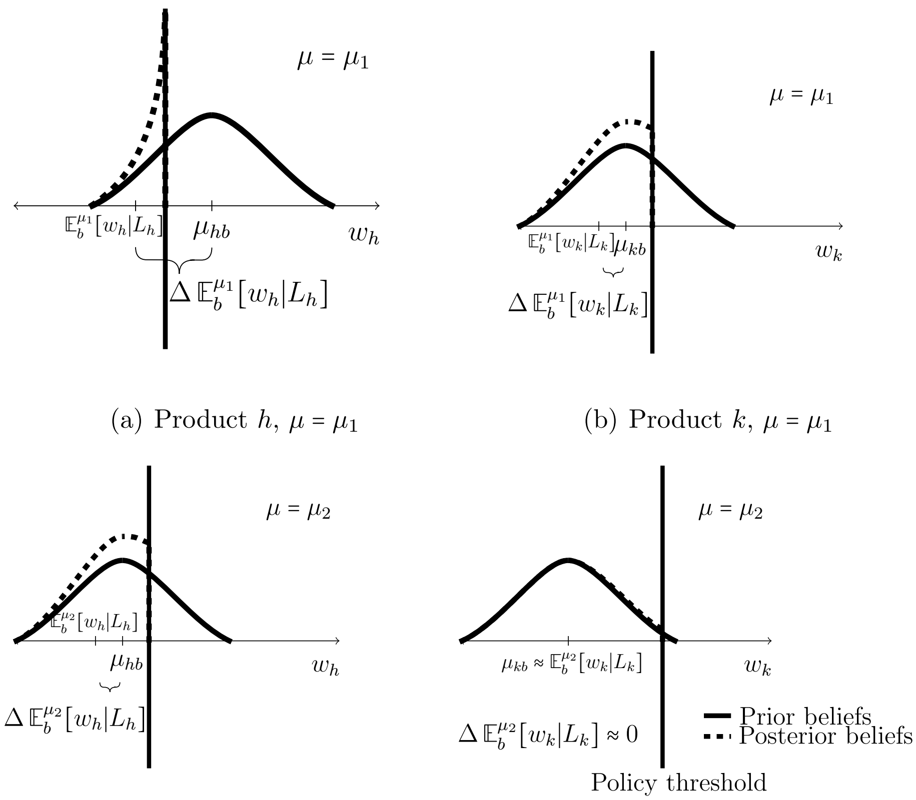

(a) Product h, $\mu = \mu_1$

(b) Product k, $\mu = \mu_1$

(c) Product h, $\mu = \mu_2$

(d) Product k, $\mu = \mu_2$

Notes: The figure illustrates the changes in beliefs about nutritional content, w, for products h and k when they do not receive a label. Product h is believed to have a higher concentration of the critical nutrient, w, than product k. Larger values of μ shift the distribution of beliefs to the right. In each panel, the yellow solid line represents the distribution of prior beliefs and the blue dashed line represents the distribution of posterior beliefs. In panels (a) and (b), we plot the distribution of prior and posterior beliefs when $\mu = \mu_1 > \mu_2$ for products h and k, respectively. In panels (c) and (d) we plot the distribution of prior and posterior beliefs when $\mu = \mu_2 < \mu_1$ for products h and k, respectively. The figure shows that changes in beliefs upon not receiving a label are larger when μ is larger. Moreover, the differences in changes in beliefs between products h and k is also larger when μ is larger.

Figure 3.3: Model-implied change in expected utility for product h and k at different values of μ upon not receiving a label

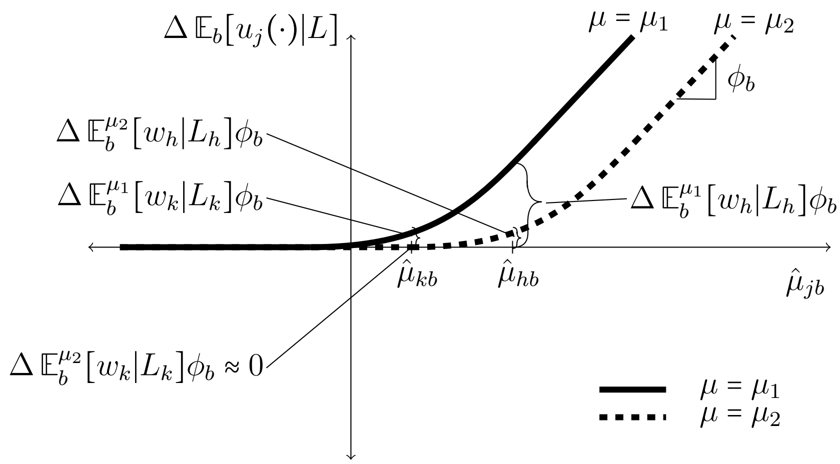

Notes: The figure illustrates the change in expected utility from consuming product j as a function of $\tilde{\mu}_{jb}$ for two different values of μ, where $\tilde{\mu}_{jb}$ is the average survey response regarding the expected value of nutritional content of product j among consumers of type b. The yellow solid line conveys this relationship for $\mu = \mu_1$ and the blue dashed line for $\mu = \mu_2$. The figure shows that different values of μ imply different changes in expected utility for products that do not get a label. Lower values of μ translate into small changes in expected utility for a broader set of products.

Figure 3.4: Willigness to pay to reduce sugar and caloric concentration for low- and high-SES consumers

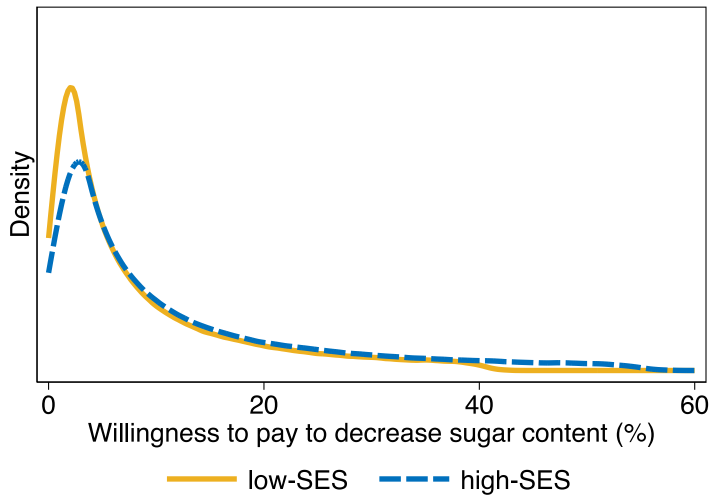

(a) Sugar

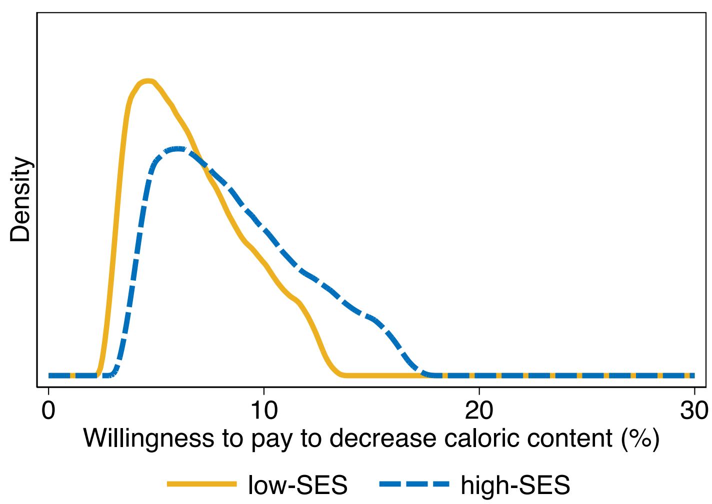

(b) Calories

Notes: The figure presents willingness to pay, as a percentage of the average price of cereal, among low- and high-SES consumers to reduce the sugar and caloric concentration of products by 1 standard deviation while keeping the taste constant. To calculate willingness to pay, we use the following formula: $\text{wtp}_i = \frac{\phi_i}{\alpha_i}\frac{sd(w_{jt})}{\bar{p}_{jt}}$. The parameters that govern the distributions of ϕ_i and α_i are reported in Table 3.1.

Figure 3.5: Predicted probability of bunching

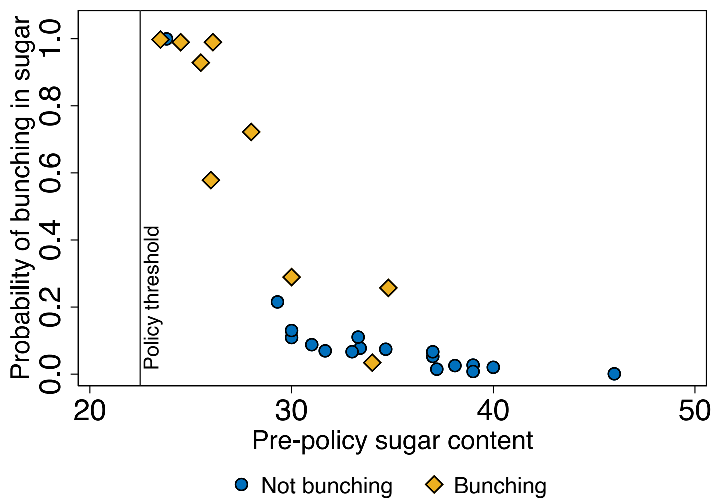

(a) Sugar

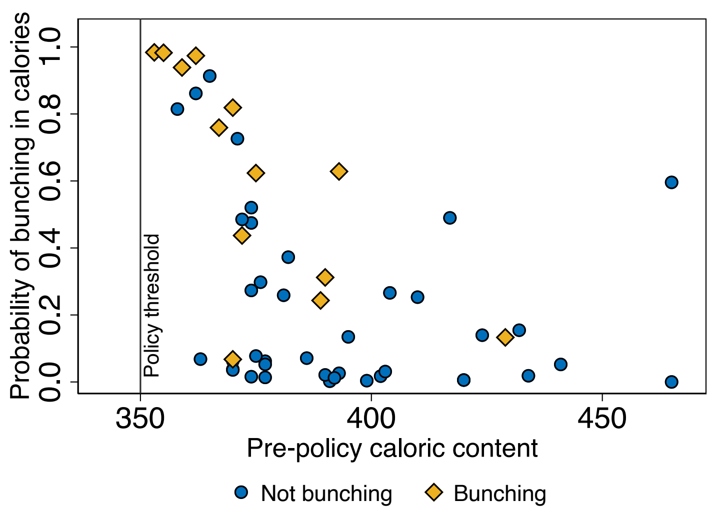

(b) Calories

Notes: The figure shows the predicted probability of each product bunching in sugar and calories as a function of the pre-policy nutritional content and the distance from the regulatory threshold for each critical nutrient. In Panel (a), we focus on sugar content. Products in yellow diamonds are products that bunched in the data and crossed the sugar policy threshold. Products in blue circles are products that did not bunch and received a "high-in-sugar" label. In Panel (b), we focus on caloric content. Products in yellow diamonds are products that bunched in the data and crossed the calorie policy threshold. Products in blue circles are products that did not bunch and received a "high-in-calorie" label.

Figure 3.6: Changes in consumer welfare under different counterfactuals

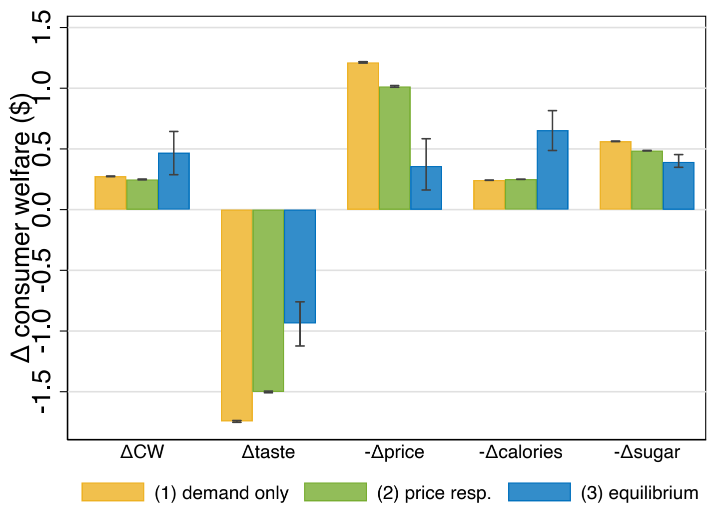

Notes: The first three bars of the figure show the changes in consumer welfare from counterfactual (0) to counterfactuals (1), (2), and (3), respectively. The remaining bars decompose these changes into changes in taste/experience of consuming cereal, changes in price paid, changes in calorie intake, and changes in sugar intake. Each bar is normalized to show the contribution of each dimension to consumer welfare in dollars. For example, a positive value for the contribution of caloric intake means that consumers are consuming lower quantities of calories under that counterfactual. We present 90% confidence intervals from the Monte Carlo simulations. Counterfactual (3) has larger confidence intervals due to variation in Λ_j that does not show up when firms do not reformulate products.

Figure 3.7: Changes in consumer welfare under food labels and sugar taxes

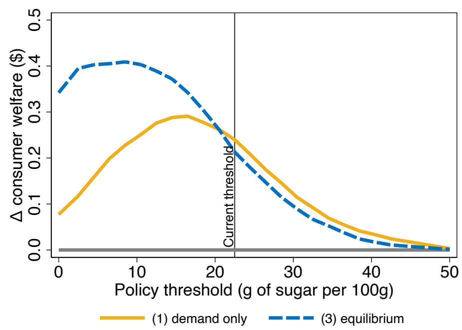

(a) Food labels

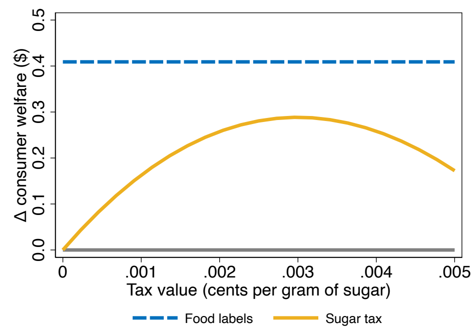

(b) Sugar taxes

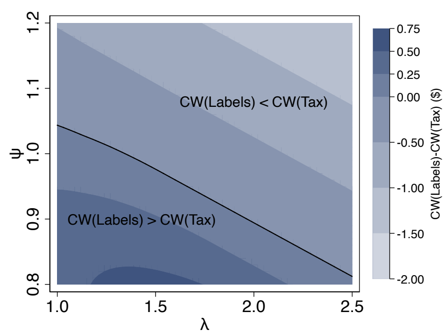

(c) Policy comparison

Notes: Panels (a) and (b) plot the average change in consumer welfare under counterfactuals (1) and (3) relative to counterfactual (0). Panel (a) shows the gains in consumer welfare under a food labeling policy at different regulatory thresholds, and panel (b) shows the gains in consumer welfare under different tax values. Panel (c) shows a contour plot that represents the difference in gains in consumer welfare between a food labeling policy and sugar taxes as a function of λ, the parameter that accounts for additional market imperfections, and ψ, the marginal value of public funds under counterfactual (3). For each value of λ and ψ, we choose policy thresholds and tax values that maximize consumer welfare. In the bottom-left side of the box, consumer welfare gains under a food labeling policy is larger than under optimal sugar taxes (CW(Labels) > CW(Tax)). In the upper-right side of the box, consumer welfare gains under a food labeling policy is smaller than under optimal sugar taxes (CW(Labels) < CW(Tax)).

Tables

Table 3.1: Estimated demand parameters

Panel A: Preferences for price and healthiness (α_i, ϕ_i)								
	First moments				Second moments			
	low-SES		high-SES		low-SES		high-SES	
Price (α_i)	$\bar{\alpha}_l$	0.255***	$\bar{\alpha}_h$	0.189***	σ_{α_l}	0.152***	σ_{α_h}	0.113***
		(0.072)		(0.059)		(0.034)		(0.036)
Sugar (ϕ_i^s)	$\bar{\phi}_l^s$	0.013***	$\bar{\phi}_h^s$	0.013**	$\sigma_{\phi_l^s}$	0.054	$\sigma_{\phi_h^s}$	0.055
		(0.004)		(0.005)		(0.151)		(0.153)
Calories (ϕ_i^c)	$\bar{\phi}_l^c$	0.026***	$\bar{\phi}_h^c$	0.025***	$\sigma_{\phi_l^c}$	0.028	$\sigma_{\phi_h^c}$	0.028
		(0.007)		(0.008)		(0.019)		(0.017)
Panel B: Individual preferences for different subcategories (Σ_β)								
Plain		Sugary		Chocolate		Granola		Oatmeal
$\sigma_{\beta_{r_1}}$	0.058	$\sigma_{\beta_{r_2}}$	0.195	$\sigma_{\beta_{r_3}}$	0.215	$\sigma_{\beta_{r_4}}$	0.036	$\sigma_{\beta_{r_5}}$ 0.295
	(0.145)		(0.186)		(0.139)		(0.167)	(0.361)
Panel C: Remaining parameters (ρ, μ)								
Nest parameter	ρ	0.959***						
		(0.004)						
Beliefs shifter	μ	-0.129***						
		(0.019)						

Notes: Nutritional content is measured in grams of sugar and kilocalories per gram of cereal, and prices in dollars per 100 grams of cereal. Subscripts l and h correspond to parameters for low- and high-SES consumers, respectively. For random parameters $x_i \in \{\alpha_i, \phi_i, \beta_i\}$, we report their average $\bar{x}$ and standard deviation σ_x. Standard errors are calculated using the delta method and reported in parentheses.

Table 3.2: Estimated supply parameters

Panel A: Costs to reformulate sugar					
$\bar{\mu}^s_\lambda$	-1.832**	σ^s_λ	1.143*		
	(0.839)		(0.677)		
Panel B: Costs to reformulate calories					
$\bar{\mu}^c_\lambda$	-2.349	σ^c_λ	1.874 *	ϑ^c_λ	1.546**
	(1.946)		(0.967)		(0.687)

Notes: The table presents the estimated parameters that govern the distribution of $\Lambda_j = \begin{bmatrix} \lambda^s_j & 0 \\ 0 & \lambda^c_j \end{bmatrix}$, the cost of reformulating sugar and calories. We assume that λ^n_j is drawn from a lognormal distribution with parameters $(\mu^n_\lambda, \sigma^n_\lambda)$, where $\mu^c_\lambda = \bar{\mu}^c_\lambda + \vartheta^c_\lambda \nu^s_j$ while $\mu^s_\lambda = \bar{\mu}^s_\lambda$. To estimate the parameters, we measure nutritional content in 10 grams of sugar and 100 kilocalories per 100 grams of cereal, respectively. Standard errors are presented in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Bibliography

Jason Abaluck. What Would We Eat if We Knew More: The Implications of a Large-Scale Change in Nutrition Labeling. *Working Paper*, 2011.

Jason Abaluck and Jonathan Gruber. Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. *American Economic Review*, 101(4): 1180–1210, 2011.

John M Abowd, Francis Kramarz, and David N Margolis. High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333, 1999.

Daron Acemoglu, Alex He, and Daniel le Maire. Eclipse of Rent-Sharing: The Effects of Managers' Business Education on Wages and the Labor Share in the US and Denmark. Working Paper 29874, National Bureau of Economic Research, 2022. URL http://www.nber.org/papers/w29874.

Daniel A Ackerberg and Gregory S Crawford. Estimating price elasticities in differentiated product demand models with endogenous characteristics. *Working Paper*, 2009.

Brian Adams and Kevin R. Williams. Zone Pricing in Retail Oligopoly. *American Economic Journal: Microeconomics*, 11(1):124–56, February 2019. doi: 10.1257/mic.20170130. URL http://www.aeaweb.org/articles?id=10.1257/mic.20170130.

Arturo Aguilar, Emilio Gutierrez, and Enrique Seira. The Effectiveness of Sin Food Taxes: Evidence from Mexico. *Journal of Health Economics*, 77:102455, 2021.

AHRQ. Prevention Quality Indicators Technical Specifications. https://www.qualityindicators.ahrq.gov/Modules/PQI_TechSpec_ICD10_v2019.aspx, 2019. Accessed: 2020-06-26.

Jorge Alé-Chilet. Gradually Rebuilding a Relationship: The Emergence of Collusion in Retail Pharmacies in Chile. Manuscript, 2018.

Jorge Alé-Chilet and Sarah Moshary. Beyond Consumer Switching: Supply Responses to Food Packaging and Advertising Regulations. *Marketing Science*, 41(2):243–270, 2022. doi: 10.1287/mksc.2021.1315. URL https://doi.org/10.1287/mksc.2021.1315.

Olivier Allais, Fabrice Etilé, and Sébastien Lecocq. Mandatory Labels, Taxes and Market Forces: An Empirical Evaluation of Fat Policies. *Journal of Health Economics*, 43:27–44, 2015.

Hunt Allcott. The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3):30–66, 2013.

Hunt Allcott and Christopher Knittel. Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments. *American Economic Journal: Economic Policy*, 11(1): 1–37, 2019.

Hunt Allcott, Rebecca Diamond, Jean-Pierre Dubé, Jessie Handbury, Ilya Rahkovsky, and Molly Schnell. Food Deserts and the Causes of Nutritional Inequality. *The Quarterly Journal of Economics*, 134(4):1793–1844, 2019a.

Hunt Allcott, Benjamin B. Lockwood, and Dmitry Taubinsky. Should we tax sugar-sweetened beverages? an overview of theory and evidence. *Journal of Economic Perspectives*, 33(3):202–27, 2019b. doi: 10.1257/jep.33.3.202. URL http://www.aeaweb.org/articles?id=10.1257/jep.33.3.202.

Hunt Allcott, Benjamin B Lockwood, and Dmitry Taubinsky. Regressive Sin Taxes, With an Application to the Optimal Soda Tax. *The Quarterly Journal of Economics*, 134(3): 1557–1626, 2019c.

M. J. Andrews, L. Gill, T. Schank, and R. Upward. High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias? *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 171(3):673–697, 2008.

Araucanía Cuenta. Se encuentran abiertas las inscripciones para farmacia municipal de villarrica. https://rb.gy/mhzisl, 2016. Accessed: 2021-03-25.

Sebastián Araya, Andrés Elberg, Carlos Noton, and Daniel Schwartz. Identifying Food Labeling Effects on Consumer Behavior. *Marketing Science*, 41(5):982–1003, 2022. doi: 10.1287/mksc.2022.1356. URL https://doi.org/10.1287/mksc.2022.1356.

Peter Arcidiacono, Paul B. Ellickson, Carl F. Mela, and John D. Singleton. The Competitive Effects of Entry: Evidence from Supercenter Expansion. *American Economic Journal: Applied Economics*, 12(3):175–206, July 2020. doi: 10.1257/app.20180047. URL https://www.aeaweb.org/articles?id=10.1257/app.20180047.

Arlene S Ash, Stephen F Fienberg, Thomas A Louis, Sharon-Lise T Normand, Therese A Stukel, and Jessica Utts. Statistical issues in assessing hospital performance. 2012.

Healthcare Financial Management Association et al. Achieving Operating Room Efficiency Through Process Integration. *Healthcare financial management: journal of the Healthcare Financial Management Association*, 57(3):1–112, 2003.

Juan Pablo Atal, José Igancio Cuesta, Felipe González, and Cristóbal Otero. The Economics of the Public Option: Evidence from Local Pharmaceutical Markets. Working paper, 2022a.

Juan Pablo Atal, José Ignacio Cuesta, and Morten Sæthre. Quality Regulation and Competition: Evidence from Pharmaceutical Markets. Working Paper 30325, National Bureau of Economic Research, August 2022b.

David Atkin, Benjamin Faber, and Marco Gonzalez-Navarro. Retail Globalization and Household Welfare: Evidence from Mexico. *Journal of Political Economy*, 126(1):1–73, 2018.

Anthony Atkinson and Joseph E Stiglitz. *Lectures on Public Economics*. McGraw-Hill, 1980.

Michel Azulai, Imran Rasul, Daniel Rogger, and Martin Williams. Can Training Improve Organizational Culture? Experimental Evidence from Ghana's Civil Service. Working paper, 2020.

Matthew Backus, Christopher Conlon, and Michael Sinkinson. Common Ownership and Competition in the Ready-to-Eat Cereal Industry. *Working Paper*, 2021.

Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. In *Rationality and Equilibrium*, pages 217–241. Springer, 2006.

Katherine Baicker and Amy Finkelstein. The Impact of Medicaid Expansion on Voter Participation: Evidence from the Oregon Health Insurance Experiment. *Quarterly Journal of Political Science*, 14(4):383–400, 2019.

Andrew C. Baker, David F. Larcker, and Charles C.Y. Wang. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2):370–395, 2022. ISSN 0304-405X. doi: https://doi.org/10.1016/j.jfineco.2022.01.004. URL https://www.sciencedirect.com/science/article/pii/S0304405X22000204.

Abhijit Banerjee, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto. Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia. *Journal of Political Economy*, 127(1):101–137, 2019.

Nano Barahona, Joshua Kim, Cristóbal Otero, and Sebastián Otero. On the design of food labeling policies. *Working Paper*, 2022.

Eduardo Barros, Alejandro Weber, and Daniel Díaz. Convenios de Desempeño en la Alta Dirección Pública. Orientaciones de Optimización como Herramienta de Destión del Desempeño. In Isabel Aninat and Slaven Razmilic, editors, *Un Estado para la Ciudadanía. Estudios para su modernización*. Centro de Estudios Públicos, CEP, 2018.

Emek Basker. The Causes and Consequences of Wal-Mart's Growth. *Journal of Economic Perspectives*, 21(3):177–198, 2007.

P Beato and A. Mas-Colell. The Marginal Cost Pricing as a Regulation Mechanism in Mixed Markets. *The Performance of Public Enterprises, Amsterdam: North-Holland*, 1984.

Stefan Bender, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter. Management Practices, Workforce Selection, and Productivity. *Journal of Labor Economics*, 36(S1):S371–S409, 2018. doi: 10.1086/694107. URL https://doi.org/10.1086/694107.

Alejandra Benítez, Andrés Hernando, and Carolina Velasco. Encuesta de Presupuestos Familiares y Gasto en Salud: Una Primera Mirada. *Centro de Estudios Públicos, Punto de Referencia, N. 484*, 2018.

Morten Bennedsen, Francisco Pérez-González, and Daniel Wolfenzon. Do CEOs matter? Evidence from hospitalization events. *The Journal of Finance*, 75(4):1877–1911, 2020.

Daniel Bennett and Wesley Yin. The Market for High-Quality Medicine: Retail Chain Entry and Drug Quality in India. *Review of Economics and Statistics*, 101(1):76–90, March 2019. URL https://ideas.repec.org/a/tpr/restat/v101y2019i1p76-90.html.

Lauren Falcao Bergquist and Michael Dinerstein. Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya. *American Economic Review*, 110(12):3705–47, December 2020. doi: 10.1257/aer.20171397. URL https://www.aeaweb.org/articles?id=10.1257/aer.20171397.

B. Douglas Bernheim and Dmitry Taubinsky. Behavioral public economics. *In: Handbook of Behavioral Economics—Foundations and Applications*, 1(5):381–516, 2018. URL http://www.jstor.org/stable/2117898.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890, 1995. URL https://ideas.repec.org/a/ecm/emetrp/v63y1995i4p841-90.html.

Marianne Bertrand. Coase Lecture – The Glass Ceiling. *Economica*, 85(338):205–231, 2018. doi: 10.1111/ecca.12264. URL https://ideas.repec.org/a/bla/econom/v85y2018i338p205-231.html.

Marianne Bertrand and Antoinette Schoar. Managing with Style: The Effect of Managers on Firm Policies. *The Quarterly Journal of Economics*, 118(4):1169–1208, 2003. ISSN 0033-5533. doi: 10.1162/003355303322552775. URL https://doi.org/10.1162/003355303322552775.

Marianne Bertrand, Robin Burgess, Arunish Chawla, and Guo Xu. The Glittering Prizes: Career Incentives and Bureaucrat Performance. *The Review of Economic Studies*, 87(2): 626–655, 2020. ISSN 0034-6527. doi: 10.1093/restud/rdz029. URL https://doi.org/10.1093/restud/rdz029.

Timothy Besley and Anne Case. Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits. *Quarterly Journal of Economics*, 110(3):769–798, 1995.

Timothy Besley, Robin Burgess, Adnan Khan, and Guo Xu. Bureaucracy and Development. *Annual Review of Economics*, 14(1):null, 2022. doi: 10.1146/annurev-economics-080521-011950. URL https://doi.org/10.1146/annurev-economics-080521-011950.

Michael Carlos Best, Jonas Hjort, and David Szakonyi. Individuals and Organizations as Sources of State Effectiveness. 2022.

Barbara Biasi. The Labor Market for Teachers under Different Pay Schemes. *American Economic Journal: Economic Policy*, 13(3):63–102, August 2021. doi: 10.1257/pol.20200295. URL https://www.aeaweb.org/articles?id=10.1257/pol.20200295.

Nicholas Bloom, Renata Lemos, Raffaella Sadun, and John Van Reenen. Does Management Matter in Schools? *The Economic Journal*, 125(584):647–674, 2015a. doi: 10.1111/ecoj.12267. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12267.

Nicholas Bloom, Carol Propper, Stephan Seiler, and John Van Reenen. The Impact of Competition on Management Quality: Evidence from Public Hospitals. *The Review of Economic Studies*, 82(2):457–489, 2015b. ISSN 0034-6527. doi: 10.1093/restud/rdu045. URL https://doi.org/10.1093/restud/rdu045.

Nicholas Bloom, Renata Lemos, Raffaella Sadun, and John Van Reenen. Healthy Business? Managerial Education and Management in Health Care. *The Review of Economics and Statistics*, 102(3):506–517, 07 2020. ISSN 0034-6535. doi: 10.1162/rest_a_00847. URL https://doi.org/10.1162/rest_a_00847.

B Bollinger, P Leslie, and A Sorensen. Calorie Posting in Chain Restaurants. *American Economic Journal: Economic Policy*, 3:91–128, 2011. ISSN 1945-7731. doi: 10.3386/w15648. URL http://dx.doi.org/10.3386/w15648.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting Event Study Designs: Robust and Efficient Estimation, 2022a. URL https://arxiv.org/abs/2108.12419.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting Event Study Designs: Robust and Efficient Estimation. Mimeo, 2022b.

Javier Brugués. The Effects of Public Procurement on Medicine Supply. Manuscript, 2020.

Matias Busso and Sebastian Galiani. The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, 11(1):33–56, January 2019. doi: 10.1257/app.20150310. URL http://www.aeaweb.org/articles?id=10.1257/app.20150310.

CADP. Estado del Sistema de Alta Dirección Pública al 2016. Rendición de Cuentas a las Comisiones de Hacienda del Congreso Nacional. Technical report, Consejo de Alta Dirección Pública, 2017.

David Card, Jörg Heining, and Patrick Kline. Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013. ISSN 0033-5533. doi: 10.1093/qje/qjt006. URL https://doi.org/10.1093/qje/qjt006.

CASEN. Síntesis de Resultados CASEN 2017: Salud, 2017. URL https://www.minsal.cl/wp-content/uploads/2018/10/CASEN-Salud-2017.pdf.

Richard E. Caves, Michael D. Whinston, Mark A. Hurwitz, Ariel Pakes, and Peter Temin. Patent Expiration, Entry, and Competition in the U.S. Pharmaceutical Industry. *Brookings Papers on Economic Activity. Microeconomics*, 1991:1–66, 1991. ISSN 10578641. URL http://www.jstor.org/stable/2534790.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454, 05 2019a. ISSN 0033-5533. doi: 10.1093/qje/qjz014. URL https://doi.org/10.1093/qje/qjz014.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics*, 134(3):1405–1454, 05 2019b.

David Jr. Chan, David Card, and Lowell Taylor. Is There a VA Advantage? Evidence from Dually Eligible Veterans. Working Paper 29765, National Bureau of Economic Research, 2022. URL https://www.nber.org/papers/w29765.

Amitabh Chandra, Amy Finkelstein, Adam Sacarny, and Chad Syverson. Health care exceptionalism? performance and allocation in the us health care sector. *American Economic Review*, 106(8):2110–44, August 2016. doi: 10.1257/aer.20151080. URL https://www.aeaweb.org/articles?id=10.1257/aer.20151080.

Yongmin Chen and Michael H. Riordan. Price-Increasing Competition. *RAND Journal of Economics*, 39(4):1042–1058, 2008.

Alberto Chong, Ana De La O, Dean Karlan, and Leonard Wantchekon. Does Corruption Information Inspire the Fight or Quash the Hope? A Field Experiment in Mexico on Voter Turnout, Choice, and Party Identification. *Journal of Politics*, 77(1):55–71, 2015.

Prithwiraj Choudhury, Tarun Khanna, and Christos A Makridis. Do Managers Matter? A Natural Experiment from 42 R
amp;D Labs in India. *The Journal of Law, Economics, and Organization*, 36(1):47–83, 12 2019. ISSN 8756-6222. doi: 10.1093/jleo/ewz019. URL https://doi.org/10.1093/jleo/ewz019.

Clave9. Municipio abrió la farmacia municipal de angol. https://rb.gy/uwpzfz, 2017. Accessed: 2021-03-25.

Joshua D. Clinton and Michael W. Sances. The Politics of Policy: The Initial Mass Political Effects of Medicaid Expansion in the States. *American Political Science Review*, 112(1): 167?185, 2018. doi: 10.1017/S0003055417000430.

CNEP. Uso Eficiente de Quirófanos Electivos y Gestión de Lista de Espera Quirúrgica No GES. Technical report, Comisión Nacional de Evaluación y Productividad, 2020.

Emanuele Colonnelli, Mounu Prem, and Edoardo Teso. Patronage and Selection in Public Sector Organizations. *American Economic Review*, 110(10):3071–99, 2020.

Teresa Correa, Marcela Reyes, Lindsey Smith Taillie, Camila Corvalán, and Francesca R Dillman Carpentier. Food advertising on television before and after a national unhealthy food marketing regulation in chile, 2016–2017. *American Journal of Public Health*, 110 (7):1054–1059, 2020.

CPPUC. Informe Final: Convenios de Desempeño. Rediseño de los convenios de desempeño de los altos directivos públicos. Technical report, Centro UC Políticas Públicas, 2013.

Helmuth Cremer, Maurice Marchand, and Jacques-François Thisse. Mixed Oligopoly with Differentiated Products. *International Journal of Industrial Organization*, 9(1):43–53, 1991.

Vilsa Curto, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya. Healthcare Spending and Utilization in Public and Private Medicare. *American Economic Journal: Applied Economics*, 11(2):302–332, 2019.

David M. Cutler and Wendy Everett. Thinking Outside the Pillbox: Medication Adherence as a Priority for Health Care Reform. *New England Journal of Medicine*, 362(17):1553–1555, 2010.

Ricardo Dahis, Laura Schiavon, and Thiago Scot. Selecting Top Bureaucrats: Admission Exams and Performance in Brazil. *Review of Economics and Statistics*, page Forthcoming, 2022.

Ernesto Dal Bó, Frederico Finan, and Martín A. Rossi. Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *The Quarterly Journal of Economics*, 128(3):1169–1218, 2013. ISSN 0033-5533. doi: 10.1093/qje/qjt008. URL https://doi.org/10.1093/qje/qjt008.

Ernesto Dal Bó, Pedro Dal Bó, and Erik Eyster. The Demand for Bad Policy When Voters Underappreciate Equilibrium Effects. *Review of Economic Studies*, 85(2):964–998, 2018.

Clément De Chaisemartin and Xavier d'Haultfoeuille. Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–96, 2020.

Giovanni De Fraja and Flavio Delbono. Alternative strategies of a public enterprise in oligopoly. *Oxford Economic Papers*, 41(2):302–311, 1989.

Decreto 1671 Exento. Aprueba Norma Técnica que Establece Uso de Formularo Informe Estadístico de Egreso Hospitalario para la Producción de Información Estadística Sobre Causas de Egreso Hospitalario y Variables Asociadas. Available at https://www.bcn.cl/leychile/navegar?i=1019779, 2010. *Ministerio de Salud*. Accessed: 2022-02-09.

Decreto 38. Reglamento Orgánico de los Establecimientos de Salud de Menor Complejidad y de los Establecimientos de Autogestión en Red. Available at: https://www.bcn.cl/leychile/navegar?i=245619, 2005. Accessed: 2022-07-22.

DEIS. Egresos Hospitalarios. Available at http://www.deis.cl/estadisticas-egresoshospitalarios/, 2019. Accessed: 2020-02-09.

Stefano DellaVigna and Matthew Gentzkow. Uniform Pricing in U.S. Retail Chains. *Quarterly Journal of Economics*, 134(4):2011–2084, 06 2019.

Brian Denton, James Viapiano, and Andrea Vogl. Optimization of Surgery Sequencing and Scheduling Decisions Under Uncertainty. *Health Care Management Science*, 10(1): 13–24, 2007. doi: 10.1007/s10729-006-9005-4. URL https://doi.org/10.1007/s10729-006-9005-4.

Erika Deserranno, Stefano Caria, Philipp Kastrau, and Gianmarco León-Ciliotta. The Allocation of Incentives in Multi-Layered Organizations. Economics Working Papers 1838, Department of Economics and Business, Universitat Pompeu Fabra, 2022. URL https://ideas.repec.org/p/upf/upfgen/1838.html.

Diario Concepción. Hualpén inauguró farmacia y óptima municipal. rb.gy/jtkbid, 2017. Accessed: 2021-03-25.

Diario Financiero. Farmacias populares son efectivas en bajar los precios de los medicamentos básicos. shorturl.at/rsBO6, 2022. Accessed: 2022-10-14.

Marias Dias and Claudio Ferraz. Voting for Quality? The Impact of School Performance Information on Electoral Outcomes. PUC-Rio Working Paper 668, 2019.

Michael Dinerstein and Troy Smith. Quantifying the Supply Response of Private Schools to Public Policies. *American Economic Review*, 111(10):3376–3417, 2021.

Michael Dinerstein, Christopher Neilson, and Sebastián Otero. The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy. Working Paper, 2022.

Joseph Doyle, John Graves, and Jonathan Gruber. Evaluating Measures of Hospital Quality: Evidence from Ambulance Referral Patterns. *The Review of Economics and Statistics*, 101(5):841–852, 12 2019. ISSN 0034-6535. doi: 10.1162/rest_a_00804. URL https://doi.org/10.1162/rest_a_00804.

Joseph J. Doyle, John A. Graves, Jonathan Gruber, and Samuel A. Kleiner. Measuring Returns to Hospital Care: Evidence from Ambulance Referral Patterns. *Journal of Political Economy*, 123(1):170–214, 2015. doi: 10.1086/677756. URL https://doi.org/10.1086/677756.

Michaela Draganska, Michael Mazzeo, and Katja Seim. Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions. *QME*, 7(2):105–146, 2009.

David Dranove and Ginger Zhe Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–63, 2010.

David Dranove, Daniel Kessler, Mark McClellan, and Mark Satterthwaite. Is More Information Better? The Effects of Report Cards on Health Care Providers. *Journal of Political Economy*, 111(3):555–588, 2003.

Marco Duarte, Lorenzo Magnolfi, and Camilla Roncoroni. The Competitive Conduct of Consumer Cooperatives. *Working Paper*, 2021.

Pierre Dubois and Laura Lasio. Identifying Industry Margins with Price Constraints: Structural Estimation on Pharmaceuticals. *American Economic Review*, 108(12):3685–3724, December 2018. doi: 10.1257/aer.20140202. URL http://www.aeaweb.org/articles?id=10.1257/aer.20140202.

Pierre Dubois, Rachel Griffith, and Martin O'Connell. The Effects of Banning Advertising in Junk Food Markets. *The Review of Economic Studies*, 85(1):396–436, 04 2017. ISSN 0034-6527. doi: 10.1093/restud/rdx025. URL https://doi.org/10.1093/restud/rdx025.

Pierre Dubois, Rachel Griffith, and Martin O'Connell. How well targeted are soda taxes? *American Economic Review*, 110(11):3661–3704, November 2020. doi: 10.

1257/aer.20171898. URL https://www.aeaweb.org/articles?id=10.1257/aer.20171898.

Pierre Dubois, Yassine Lefouilli, and Stéphane Straub. Pooled Procurement of Drugs in Low and Middle Income Countries. *European Economic Review*, 132:103655, 2021.

Pierre Dubois, Ashvin Gandhi, and Shoshana Vasserman. Bargaining and International Reference Pricing in the Pharmaceutical Industry. *Working Paper*, 2022.

Mark Duggan and Fiona M. Scott Morton. The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing. *Quarterly Journal of Economics*, 121(1):1–30, 2006. doi: 10.1093/qje/121.1.1. URL http://dx.doi.org/10.1093/qje/121.1.1.

Mark Duggan, Atul Gupta, Emilie Jackson, and Zachary Templeton. The Impact of Privatization: Evidence from the Hospital Sector. Working paper, 2022.

Guillermo Durán, Pablo A. Rey, and Patricio Wolff. Solving the Operating Room Scheduling Problem with Prioritized Lists of Patients. *Annals of Operations Research*, 258(2): 395–414, 2017. doi: 10.1007/s10479-016-2172-x. URL https://doi.org/10.1007/s10479-016-2172-x.

El Austral. Hoy comenzó a funcionar la farmacia popular de padre las casas. https://rb.gy/91whig, 2017. Accessed: 2021-03-25.

Anne Elixhauser, Claudia Steiner, D Robert Harris, and Rosanna M Coffey. Comorbidity Measures for Use with Administrative Data. *Medical care*, pages 8–27, 1998.

Dennis Epple and Richard E Romano. Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects. *American Economic Review*, pages 33–62, 1998.

Jennifer Falbe, Nadia Rojas, Anna H Grummon, and Kristine A Madsen. Higher Retail Prices of Sugar-Sweetened Beverages 3 Months after Implementation of an Excise Tax in Berkeley, California. *American Journal of Public Health*, 105(11):2194–2201, 2015.

Jennifer Falbe, Hannah R Thompson, Christina M Becker, Nadia Rojas, Charles E McCulloch, and Kristine A Madsen. Impact of the Berkeley Excise Tax on Sugar-Sweetened Beverage Consumption. *American Journal of Public Health*, 106(10):1865–1871, 2016.

Ying Fan. Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review*, 103(5):1598–1628, 2013.

C Edward Fee, Charles J Hadlock, and Joshua R Pierce. Managers with and without Style: Evidence Using Exogenous Variation. *The Review of Financial Studies*, 26(3):567–601, 2013.

Alessandra Fenizia. Managers and Productivity in the Public Sector. *Econometrica*, 90(3):1063–1084, 2022. doi: https://doi.org/10.3982/ECTA19244. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA19244.

Claudio Ferraz and Frederico Finan. Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes. *Quarterly journal of economics*, 123(2): 703–745, 2008.

Frederico Finan, Ben Olken, and Rohini Pande. The Personnel Economics of the State. In Abhijit Banerjee and Esther Duflo, editors, *Handbook of Field Experiments*. North Holland, 2017.

Amy Finkelstein, Matthew Gentzkow, Peter Hull, and Heidi Williams. Adjusting Risk Adjustment — Accounting for Variation in Diagnostic Intensity. *New England Journal of Medicine*, 376(7):608–610, 2017. doi: 10.1056/NEJMp1613238. URL https://doi.org/10.1056/NEJMp1613238. PMID: 28199802.

Eric A. Finkelstein, Kiersten L. Strombotne, Nadine L. Chan, and James Krieger. Mandatory Menu Labeling in One Fast-Food Chain in King County, Washington. *American Journal of Preventive Medicine*, 40(2):122–127, 2011. ISSN 07493797. doi: 10.1016/j.amepre.2010.10.019. URL http://dx.doi.org/10.1016/j.amepre.2010.10.019.

FNE. Denuncia contra Carozzi por Conductas Anticompetitivas. Technical report, Fiscalía Nacional Económica, Santiago, Chile, 2014. URL https://www.fne.gob.cl/wp-content/uploads/2014/06/inpu_018_2014.pdf.

FNE. Estudio de Mercado sobre Medicamentos. *Fiscalía Nacional Económica*, 2019.

Julia Fonseca and Adrien Matray. Financial Inclusion, Economic Development, and Inequality: Evidence from Brazil. Manuscript, 2022.

Richard G. Frank and David S. Salkever. Generic Entry and the Pricing of Pharmaceuticals. *Journal of Economics and Management Strategy*, 6(1):75–90, 1997. ISSN 00384038. URL http://www.jstor.org/stable/1060523.

Carola Frydman. Rising Through the Ranks: The Evolution of the Market for Corporate Executives, 1936–2003. *Management Science*, 65(11):4951–4979, 2019. doi: 10.1287/mnsc.2018.3080. URL https://doi.org/10.1287/mnsc.2018.3080.

Martin Gaynor, Rodrigo Moreno-Serra, and Carol Propper. Death by Market Power: Reform, Competition, and Patient Outcomes in the National Health Service. *American Economic Journal: Economic Policy*, 5(4):134–66, 2013. doi: 10.1257/pol.5.4.134. URL https://www.aeaweb.org/articles?id=10.1257/pol.5.4.134.

Alan Gerber, James Gimpel, Donald Green, and Daron Shaw. How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment. *American Political Science Review*, 105(1):135–150, 2011.

Amory Gethin. Revisiting Global Poverty Reduction: Public-Private Complementarities and the Rise of Public Goods. Manuscript, 2022.

John Geweke, Gautam Gowrisankaran, and Robert J. Town. Bayesian Inference for Hospital Quality in a Selection Model. *Econometrica*, 71(4):1215–1238, 2003. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1555495.

Greer K Gosnell, John A List, and Robert D Metcalfe. The Impact of Management Practices on Employee Productivity: A Field Experiment with Airline Captains. *Journal of Political Economy*, 128(4):1195–1233, 2020.

Gautam Gowrisankaran, Aviv Nevo, and Robert Town. Mergers When Prices are Negotiated: Evidence from the Hospital Industry. *American Economic Review*, 105(1):172–203, 2015.

Michael Greenstone, Paul Oyer, and Annette Vissing-Jorgensen. Mandated Disclosure, Stock Returns, and the 1964 Securities Acts Amendments. *The Quarterly Journal of Economics*, 121(2):399–460, 2006.

Francesca Guerriero and Rosita Guido. Operational Research in the Management of the Operating Theatre: A Survey. *Health Care Management Science*, 14(1):89–114, 2011. doi: 10.1007/s10729-010-9143-6. URL https://doi.org/10.1007/s10729-010-9143-6.

Atul Gupta. Impacts of performance pay for hospitals: The readmissions reduction program. *American Economic Review*, 111(4):1241–83, 2021. doi: 10.1257/aer.20171825. URL https://www.aeaweb.org/articles?id=10.1257/aer.20171825.

Donald C Hambrick and Phyllis A Mason. Upper echelons: The Organization as a Reflection of its Top Managers. *Academy of management review*, 9(2):193–206, 1984.

Jesse Handbury and Sarah Moshary. School Food Policy Affects Everyone: Retail Responses to the National School Lunch Program. Manuscript, 2021.

Benjamin R. Handel and Jonathan T. Kolstad. Health insurance for "humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8): 2449–2500, August 2015. doi: 10.1257/aer.20131126. URL http://www.aeaweb.org/articles?id=10.1257/aer.20131126.

Jake Haselswerdt. Expanding Medicaid, Expanding the Electorate: The Affordable Care Act's Short-Term Impact on Political Participation. *Journal of Health Politics, Policy and Law*, 42(4):667–695, 2017.

Justine S Hastings and Jeffrey M Weinstein. Information, School Choice, and Academic Achievement: Evidence from Two Experiments. *The Quarterly journal of Economics*, 123 (4):1373–1414, 2008.

Jerry Hausman. Consumer Benefits from Increased Competition in Shopping Outlets: Measuring the Effect of Wal-Mart. *Journal of Applied Econometrics*, 22(7):1157–1177, 2007.

Biyu He, Franklin Dexter, Alex Macario, and Stefanos Zenios. The Timing of Staffing Decisions in Hospital Operating Rooms: Incorporating Workload Heterogeneity into the Newsvendor Problem. *Manufacturing & Service Operations Management*, 14(1):99–114, 2012. doi: 10.1287/msom.1110.0350. URL https://doi.org/10.1287/msom.1110.0350.

Health and Social Care Information Centre. Summary Hospital-level Mortality Indicator. 2015.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006a.

Igal Hendel and Aviv Nevo. Sales and consumer inventory. *The RAND Journal of Economics*, 37(3):543–561, 2006b.

Nathaniel Hendren and Ben Sprung-Keyser. A Unified Welfare Analysis of Government Policies. *The Quarterly Journal of Economics*, 135(3):1209–1318, 2020.

Mitchell Hoffman and Steven Tadelis. People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis. *Journal of Political Economy*, 129(1):243–285, 2021.

Sébastien Houde. The Incidence of Coarse Certification: Evidence from the ENERGY STAR Program. *Working Paper*, 2018.

Caroline M Hoxby. Does Competition Among Public Schools Benefit Students and Taxpayers? *American Economic Review*, 90(5):1209–1238, 2000.

Chang-Tai Hsieh, Erik Hurst, Charles I. Jones, and Peter J. Klenow. The allocation of talent and u.s. economic growth. *Econometrica*, 87(5):1439–1474, 2019. doi: https://doi.org/10.3982/ECTA11427. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA11427.

Kilian Huber, Volker Lindenthal, and Fabian Waldinger. Discrimination, managers, and firm performance: Evidence from "aryanizations" in nazi germany. *Journal of Political Economy*, 129(9):2455–2503, 2021. doi: 10.1086/714994. URL https://doi.org/10.1086/714994.

Peter Hull. Estimating Hospital Quality with Quasi-Experimental Data. Working papers, 2020.

Gastón Illanes and Sarah Moshary. Deregulation through Direct Democracy: Lessons from Liquor Markets. Manuscript, 2020.

Pauline M. Ippolito and Alan D. Mathios. Information, advertising and health choices: A study of the cereal market. *The RAND Journal of Economics*, 21(3):459–480, 1990. ISSN 07416261. URL http://www.jstor.org/stable/2555620.

Pauline M. Ippolito and Alan D. Mathios. Information and advertising: The case of fat consumption in the united states. *American Economic Review*, 85(2):91–95, 1995. ISSN 00028282. URL http://www.jstor.org/stable/2117898.

Katharina Janke, Carol Propper, and Raffaella Sadun. The Impact of CEOs in the Public Sector: Evidence from the English NHS. Working Paper 18-075, Harvard Business School, 2020.

Panle Jia. What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry. *Econometrica*, 76(6):1263–1316, 2008.

Diego Jiménez-Hernández and Enrique Seira. Should the Government Sell You Goods? Evidence from the Milk Market in Mexico. 2022.

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.

Zi Yang Kang and Shoshana Vasserman. Robust Bounds for Welfare Analysis. Technical report, National Bureau of Economic Research, 2022.

Lawrence F. Katz. Efficiency Wage Theories: A Partial Evaluation. *NBER Macroeconomics Annual*, 1:235–276, 1986. ISSN 08893365, 15372642. URL http://www.jstor.org/stable/3585171.

Chad Kendall, Tommaso Nannicini, and Francesco Trebbi. How Do Voters Respond to Information? Evidence from a Randomized Campaign. *American Economic Review*, 105 (1):322–53, January 2015. doi: 10.1257/aer.20131063. URL http://www.aeaweb.org/articles?id=10.1257/aer.20131063.

Adnan Q. Khan, Asim I. Khwaja, and Benjamin A. Olken. Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors. *The Quarterly Journal of Economics*, 131(1):219–271, 2015. ISSN 0033-5533. doi: 10.1093/qje/qjv042. URL https://doi.org/10.1093/qje/qjv042.

Kristin Kiesel and Sofia Villas-Boas. Can Information Costs Affect Consumer Choice? Nutritional Labels in a Supermarket Experiment. *International Journal of Industrial Organization*, 31(2):153–163, 2013.

Henrik Jacobsen Kleven and Wojciech Kopczuk. Transfer Program Complexity and the Take-Up of Social Benefits. *American Economic Journal: Economic Policy*, 3(1):54–90, February 2011. doi: 10.1257/pol.3.1.54. URL https://www.aeaweb.org/articles?id=10.1257/pol.3.1.54.

Julien Lafortune, Jesse Rothstein, and Diane Whitmore Schanzenbach. School Finance Reform and the Distribution of Student Achievement. *American Economic Journal: Applied Economics*, 10(2):1–26, April 2018. doi: 10.1257/app.20160567. URL http://www.aeaweb.org/articles?id=10.1257/app.20160567.

Guillermo Latorre-Núñez, Armin Lüer-Villagra, Vladimir Marianov, Carlos Obreque, Francisco Ramis, and Liliana Neriz. Scheduling Operating Rooms with Consideration of all Resources, Post Anesthesia Beds and Emergency Surgeries. *Computers Industrial Engineering*, 97:248–257, 2016. ISSN 0360-8352. doi: https://doi.org/10.1016/j.cie.2016.05.016. URL https://www.sciencedirect.com/science/article/pii/S0360835216301577.

Timothy J Layton, Nicole Maestas, Daniel Prinz, and Boris Vabson. Private vs. Public Provision of Social Insurance: Evidence from Medicaid. Working Paper 26042, National Bureau of Economic Research, July 2019. URL http://www.nber.org/papers/w26042.

Edward P. Lazear. Performance Pay and Productivity. *American Economic Review*, 90(5): 1346–1361, December 2000. doi: 10.1257/aer.90.5.1346. URL https://www.aeaweb.org/articles?id=10.1257/aer.90.5.1346.

Kim-Anne Lê, Frédéric Robin, and Olivier Roger. Sugar replacers: from technological challenges to consequences on health. *Current Opinion in Clinical Nutrition and Metabolic Care*, 19(4):310–315, 2016.

David S. Lee. Training, Wages, and Sample selection: Estimating Sharp Bounds on Treatment Effects. *Review of Economic Studies*, 76(3):1071–1102, 2009. ISSN 00346527. doi: 10.1111/j.1467-937X.2009.00536.x.

Matthew M Lee, Jennifer Falbe, Dean Schillinger, Sanjay Basu, Charles E McCulloch, and Kristine A Madsen. Sugar-Sweetened Beverage Consumption 3 Years After the Berkeley, California, Sugar-Sweetened Beverage Tax. *American Journal of Public Health*, 109(4): 637–639, 2019.

Robin S Lee, Michael D Whinston, and Ali Yurukoglu. Structural Empirical Analysis of Contracting in Vertical Markets. In *Handbook of Industrial Organization*, volume 4, pages 673–742. Elsevier, 2021.

Ley 19,937. Modifica el D.L. N. 2,763, de 1979 con la Finalidad de Establecer una Nueva Concesión de la Autoridad Sanitaria, Distintas Modalidades de Gestión y Fortaleceer la Participación Ciudadana. Available at: https://www.bcn.cl/leychile/navegar?idNorma=221629&idVersion=2008-12-31&idParte=8721253, 2004. Accessed: 2022-07-22.

Ley 20,955. Perfecciona el Sistema de Alta Dirección Pública y Fortalece la Dirección Nacional del Servicio Civil. Available at: https://www.bcn.cl/leychile/navegar?idNorma=1095821&idParte=9741584&idVersion=2016-10-20, 2016. Accessed: 2022-07-14.

Joon Ho Lim, Rishika Rishika, Ramkumar Janakiraman, and PK Kannan. Competitive effects of front-of-package nutrition labeling adoption on nutritional quality: Evidence from facts up front–style labels. *Journal of Marketing*, 84(6):3–21, 2020.

Nicola Limodio. Bureaucrat Allocation in the Public Sector: Evidence from the World Bank. *The Economic Journal*, 131(639):3012–3040, 01 2021. ISSN 0013-0133. doi: 10.1093/ej/ueab008. URL https://doi.org/10.1093/ej/ueab008.

John A. List and Daniel M. Sturm. How Elections Matter: Theory and Evidence from Environmental Policy. *Quarterly Journal of Economics*, 121(4):1249–1281, 2006.

Alessandro Lizzeri and Nicola Persico. The Provision of Public Goods under Alternative Electoral Incentives. *American Economic Review*, 91(1):225–239, 2001.

Luca Maini and Fabio Pammolli. Reference Pricing as a Deterrent to Entry: Evidence from the European Pharmaceutical Market. 2022.

Marco Manacorda, Edward Miguel, and Andrea Vigorito. Government Transfers and Political Support. *American Economic Journal: Applied Economics*, 3(3):1–28, 2011.

David A. Matsa. Competition and Product Quality in the Supermarket Industry. *Quarterly Journal of Economics*, 126:1539–1591, 2011.

William C Merrill and Norman Schneider. Government Firms in Oligopoly Industries: A Short-run Analysis. *Quarterly Journal of Economics*, 80(3):400–412, 1966.

Robert Metcalfe, Alexandre Sollaci, and Chad Syverson. Managers and Productivity in Retail. Working paper, November 2022.

Christian Michel and Stefan Weiergraeber. Estimating industry conduct in differentiated products markets. *Working Paper*, 2018.

MINECON. Relación Entre Cantidad de Farmacias y Pobreza. *División de Estudios, Ministerio de Economía, Fomento y Turismo*, 2013.

Tatiana Mocanu. Designing Gender Equity: Evidence from Hiring Practices and Committees. Working paper, 2022.

Debi Prasad Mohapatra and Chirantan Chatterjee. Price Control and Access to Drugs: The Case of India's Malarial Market. Manuscript, 2020.

Gina S Mohr, Donald R Lichtenstein, and Chris Janiszewski. The effect of marketer-suggested serving size on consumer responses: the unintended consequences of consumer attention to calorie information. *Journal of Marketing*, 76(1):59–75, 2012.

Christine Moorman, Rosellina Ferraro, and Joel Huber. Unintended Nutrition Consequences: Firm Responses to the Nutrition Labeling and Education Act. *Marketing Science*, 31(5): 717–737, 2012.

Diana Moreira and Santiago Pérez. Civil Service Exams and Organizational Performance: Evidence from the Pendleton Act. NBER Working Papers 28665, National Bureau of Economic Research, Inc, 2022. URL https://ideas.repec.org/p/nbr/nberwo/28665.html.

Ana Moura and Pedro Pita Barros. Entry and Price Competition in the Over-the-counter Drug Market after Deregulation: Evidence from Portugal. *Health Economics*, 29(8):865–877, 2020.

Pablo Muñoz and Mounu Prem. Managers' Productivity and Recruitment in the Public Sector. 2022.

Roger B. Myerson. Moral Hazard in High Office and the Dynamics of Aristocracy. *Econometrica*, 83(6):2083–2126, 2015. doi: 10.3982/ECTA9737. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9737.

Aviv Nevo. Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, 69(2):307–342, 2001. URL https://ideas.repec.org/a/ecm/emetrp/v69y2001i2p307-42.html.

New York Times. There Should Be a Public Option for Everything. https://nyti.ms/35rJ7qi, 2019. Accessed: 2021-06-15.

Martin Nordin, Inga Persson, and Dan-Olof Rooth. Education–Occupation Mismatch: Is There an Income Penalty? *Economics of Education Review*, 29(6):1047–1059, 2010. ISSN 0272-7757. doi: https://doi.org/10.1016/j.econedurev.2010.05.005. URL https://www.sciencedirect.com/science/article/pii/S0272775710000646.

Martin O'Connell and Kate Smith. Optimal sin taxation and market power. *Working Paper*, 2021.

OECD. *Health at a Glance 2015: OECD Indicators*. OECD Publishing, Paris, 2015.

OECD. OECD Health Statistics. http://www.oecd.org/health/health-data.htm, 2022a.

OECD. Private Health Insurance Spending. Brief, March 2022b. URL https://www.oecd.org/health/Spending-on-private-health-insurance-Brief-March-2022.pdf.

Max J. Pachali, Marco J.W. Kotschedoff, Arjen van Lin, Bart J. Bronnenberg, and Erica van Herpen. How Do Nutritional Warning Labels Affect Prices? *Working Paper*, 2022.

Gerard Padró i Miquel, Nancy Qian, and Yang Yao. The Rise and Fall of Local Elections in China: Theory and Empirical Evidence on the Autocrat's Trade-off. 2018.

Daniela Pinto, Rodrigo Moreno-Serra, Gianluca Cafagna, and Laura Giles. Efficient Spending for Healthier Lives. Inter-American Development Bank, Flagship Report, 2018.

Christopher Pollitt and Geert Bouckaert. *Public Management Reform: A Comparative Analysis - Into the Age of Austerity*. Fourth edition, 2017.

Carol Propper and John Van Reenen. Can Pay Regulation Kill? Panel Data Evidence on the Effect of Labor Markets on Hospital Performance. *Journal of Political Economy*, 118 (2):222–273, 2010. doi: 10.1086/653137. URL https://doi.org/10.1086/653137.

Hude Quan, Vijaya Sundararajan, Patricia Halfon, Andrew Fong, Bernard Burnand, Jean-Christophe Luthi, L Duncan Saunders, Cynthia A Beck, Thomas E Feasby, and William A Ghali. Coding Algorithms for Defining Comorbidities in ICD-9-CM and ICD-10 Administrative Data. *Medical care*, pages 1130–1139, 2005.

Imran Rasul and Daniel Rogger. Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service. *The Economic Journal*, 128(608):413–446, 2018. doi: 10.1111/ecoj.12418. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12418.

Silvio Ravaioli. Coarse and precise information in food labeling. *Working Paper*, 2021.

Brian E Roe, Mario F Teisl, and Corin R Deans. The Economics of Voluntary versus Mandatory Labels. *Annual Review of Resources Economics*, 6(1):407–427, 2014.

Ann E. Rogers, Wei-Ting Hwang, Linda D. Scott, Linda H. Aiken, and David F. Dinges. The Working Hours Of Hospital Staff Nurses And Patient Safety. *Health Affairs*, 23 (4):202–212, 2004. doi: 10.1377/hlthaff.23.4.202. URL https://doi.org/10.1377/hlthaff.23.4.202. PMID: 15318582.

Servicio Civil. Diagnóstico de Percepciones de Altos Directivos Públicos del Sector Salud, 2014. URL https://documentos.serviciocivil.cl/actas/dnsc/documentService/downloadWs?uuid=fed6cfb1-9f6f-4318-95f7-e185348afe6e.

Andrei Shleifer. State versus Private Ownership. *Journal of Economic Perspectives*, 12(4): 133–150, 1998.

Andrei Shleifer and Robert Vishny. Politicians and Firms. *Quarterly Journal of Economics*, 109(4):995–1025, 1994.

Lynn D Silver, Shu Wen Ng, Suzanne Ryan-Ibarra, Lindsey Smith Taillie, Marta Induni, Donna R Miles, Jennifer M Poti, and Barry M Popkin. Changes in Prices, Sales, Consumer Spending, and Beverage Consumption One Year After a Tax on Sugar-Sweetened Beverages in Berkeley, California, US: A Before-and-After Study. *PLoS Medicine*, 14(4): e1002283, 2017.

Elaine Silverman and Jonathan Skinner. Medicare Upcoding and Hospital Ownership. *Journal of Health Economics*, 23(2):369–389, 2004. ISSN 0167-6296. doi: https://doi.org/10.1016/j.jhealeco.2003.09.007. URL https://www.sciencedirect.com/science/article/pii/S0167629603001206.

Yunjie Song, Jonathan Skinner, Julie Bynum, Jason Sutherland, John E. Wennberg, and Elliott S. Fisher. Regional Variations in Diagnostic Practices. *New England Journal of Medicine*, 363(1):45–53, 2010. doi: 10.1056/NEJMsa0910881. URL https://doi.org/10.1056/NEJMsa0910881. PMID: 20463332.

Amy Witkoski Stimpfel. The Longer the Shifts for Hospital Nurses, the Higher the Levels of Burnout and Patient Dissatisfaction, 2012. URL https://www.healthaffairs.org/doi/10.1377/hlthaff.2011.1377.

Lindsey Smith Taillie, Marcela Reyes, Arantxa Colchero, Barry Popkin, and Camila Corvalán. An Evaluation of Chile's Law of Food Labeling and Advertising on Sugar-Sweetened Beverage Purchases from 2015 to 2017: A Before-and-After Study. *PLoS Medicine*, 17(2): e1003015, 2020.

Rebecca Taylor, Scott Kaplan, Sofia Villas-Boas, and Kevin Jung. Soda Wars: Effect of a Soda Tax Election on Soda Purchases. *Economic Inquiry*, 57(3):1480–1496, 2019.

The Week. Why America needs a public option for mortgages. https://bit.ly/2SupUkZ, 2017. Accessed: 2021-06-15.

Christopher Timmins. Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West. *Econometrica*, 70(2):603–629, 2002.

Thomas C. Tsai, Ashish K. Jha, Atul A. Gawande, Robert S. Huckman, Nicholas Bloom, and Raffaella Sadun. Hospital Board And Management Practices Are Strongly Related To Hospital Performance On Clinical Quality Metrics. *Health Affairs*, 34(8):1304–1311, 2015. doi: 10.1377/hlthaff.2014.1282. URL https://doi.org/10.1377/hlthaff.2014.1282. PMID: 26240243.

UN. MDG Gap Task Force Report 2010. Technical report, United Nations, 2010.

Benjamin Vatter. Quality Disclosure and Regulation: Scoring Design in Medicare Advantage. *Working Paper*, 2021.

Sofia B. Villas-Boas, Kristin Kiesel, Joshua P. Berning, Hayley H. Chouinard, and Jill J. McCluskey. Consumer and strategic firm response to nutrition shelf labels. *American Journal of Agricultural Economics*, 102(2):458–479, 2020.

Joachim Voth and Guo Xu. Patronage for Productivity: Selection and Performance in the Age of Sail. 2022.

Lien Wang, Erik Demeulemeester, Nancy Vansteenkiste, and Frank E. Rademakers. Operating Room Planning and Scheduling for Outpatients and Inpatients: A Review and Future Research. *Operations Research for Health Care*, 31:100323, 2021. ISSN 2211-6923. doi: https://doi.org/10.1016/j.orhc.2021.100323. URL https://www.sciencedirect.com/science/article/pii/S2211692321000394.

Michael B. Ward, Jay P. Shimshack, Jeffrey M. Perloff, and J. Michael Harris. Effects of the Private-Label Invasion in Food Industries. *American Journal of Agricultural Economics*, 84(4):961–973, 2002. doi: https://doi.org/10.1111/1467-8276.00360. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8276.00360.

WHO. Factsheet No. 311. *World Health Organization*, 2018.

Jessica Wisdom, Julie S. Downs, and George Loewenstein. Promoting Healthy Choices: Information versus Convenience. *American Economic Journal: Applied Economics*, 2(2): 164–178, 2010. ISSN 19457782. doi: 10.1257/app.2.2.164.

Thomas G. Wollmann. Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles. *American Economic Review*, 108(6):1364–1406, June 2018. doi: 10.1257/aer.20160863. URL https://www.aeaweb.org/articles?id=10.1257/aer.20160863.

Susan E Woodward and Robert E Hall. Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence. *American Economic Review*, 102(7):3249–76, 2012.

Guo Xu. The Costs of Patronage: Evidence from the British Empire. *American Economic Review*, 108(11):3170–98, 2018. doi: 10.1257/aer.20171339. URL http://www.aeaweb.org/articles?id=10.1257/aer.20171339.

Chen Zhu, Rigoberto A. Lopez, and Xiaoou Liu. Information Cost and Consumer Choices of Healthy Foods. *American Journal of Agricultural Economics*, 98(1):41–53, 2015. ISSN 14678276. doi: 10.1093/ajae/aav057.

Appendix A

The Economics of the Public Option: Evidence from Local Pharmaceutical Markets

A.1 The determinants of retail drug prices

In the main text, we argue that two conditions that generate price differences between state-owned and private firms are the higher bargaining power of the former in the wholesale market and the exercise of market power of the latter in the retail market. In this section, we present a model that formalizes this intuition.

Setup

We consider a sequential monopoly model with Nash bargaining. An upstream monopoly produces a drug that is sold to a retail pharmacy that is a downstream monopoly. The model allows for this downstream firm to represent the private pharmacy, the public pharmacy, or some combination between them—we specify how the downstream firm's objective function captures these possibilities below. The marginal cost of the upstream monopoly is c and the wholesale price the retailer pays is t. There are no additional marginal costs downstream.

We start by introducing the objective functions of the upstream firm and the retailer. The upstream monopoly maximizes profits:

$$\Pi_U(t) = (t - c)\bar{q}(t),$$

where $\bar{q}(t) \equiv q(p(t))$ are the sales that result when the downstream retailer chooses the optimal retail price given the wholesale price t.

The downstream firm sets prices by taking into account both profits and consumer surplus, with a weight on consumer surplus equal to λ. Omitting the dependence of prices with respect to the wholesale price, the objective function of the retailer is:

$$V_D(p) = (p - t)\, q\,(p) + \lambda CS\,(p)\,,$$

where $q(p)$ is the demand function, for which we assume $q'(p) < 0$ and $q''(p) \geq 0$. The parameter λ measures the degree of alignment between the retailer and consumers. If $\lambda < 1$, the retailer values profits more than consumer welfare; $\lambda > 1$ implies that the retailer values consumer welfare more than profits; and $\lambda = 1$ means that consumer welfare and profits are valued equally by the retailer and hence that the retailer maximizes total welfare.[1] In terms of the downstream market structure, this specification of the retailer objective is akin to a mixed oligopoly model for the retail market in which private and state-owned firms compete (see, e.g., Merrill and Schneider 1966; Beato and Mas-Colell 1984; De Fraja and Delbono 1989; Cremer, Marchand, and Thisse 1991; Duarte, Magnolfi, and Roncoroni 2021).

Bargaining over wholesale price. The upstream and downstream firms bargain over wholesale prices. The wholesale price t maximizes the Nash product of the gains from trade for both firms:

$$V_D\,(p(t))^{\zeta} \times (\Pi_U\,(t))^{1-\zeta}\,,$$

where ζ is the bargaining power of the retailer.

Optimal pricing upstream and downstream. The first-order condition of the Nash bargaining problem is:

$$(t - c)q'\,(p)\,p'(t) + q = \left(\frac{\zeta}{1-\zeta}\right)\frac{t-c}{(p-t) + \lambda\frac{CS}{q}} \times q, \tag{A.1}$$

where it is useful to note that this equation simplifies to the standard first order condition of the bilateral monopoly model in the case of $\lambda = 0$, where the retailer places no weight on consumer surplus (Lee, Whinston, and Yurukoglu, 2021).

The optimal retailer price is given by:

$$p = t - \frac{q}{q'} - \lambda\frac{CS'}{q'},$$

which, by using the fact that $CS' = -q(p)$, simplifies to:

$$p = t - (1-\lambda)\frac{q}{q'},$$

which only holds when $\lambda < 1$. When $\lambda \geq 1$, the downstream firm is at a corner solution where it sets prices at marginal cost, namely $p = t$. Overall, the optimal price downstream is given

[1] See also Timmins (2002) and Gowrisankaran, Nevo, and Town (2015) for similar specifications of firm objectives when aligned with consumers.

by:

$$p = \begin{cases} t - (1 - \lambda)\frac{q}{q'} & \lambda < 1 \\ t & \lambda \geq 1. \end{cases} \tag{A.2}$$

Market outcomes are jointly determined by equations (A.1) and (A.2), and depend on the bargaining power of the retailer and the extent to which the retailer is aligned with consumers and value consumer surplus.

Comparative Statics

In this section, we deliver the main results of the model. In particular, we show how wholesale and retail prices vary with the retailer's bargaining power and market power, which depend on the parameters ζ and λ, respectively. These are the results that map to the two conditions we discuss in the main text for why public state-owned firms may offer lower prices than private firms in our setting. We start by introducing three assumptions:

Assumption 1 (Decreasing Marginal Revenue)**.** *Marginal revenue $MR(q) = p(q) + qp'(q)$ is decreasing in q, where $p(q)$ is the inverse demand curve.*

Assumption 2. $\frac{qq''}{q'^2}$ *is weakly increasing in p.*

Assumption 3. $\frac{q^2}{-q'} - CS \geq 0.$

These assumptions provide conditions under which the two comparative statics of interest hold. Assumption 1 guarantees the existence of a profit-maximizing price for a monopolist facing a convex cost function and is implied by log-concavity of demand (see e.g., Kang and Vasserman 2022). Assumption 3 is also implied by log-concavity, as shown in Section A.1. Log concavity is a commonly-used assumption in industrial organization, and hence it is not particularly restrictive (Bagnoli and Bergstrom, 2006). This property of demand ensures that the first order condition of the monopoly is sufficient for profit maximization.

We start by establishing general results for how market outcomes vary with the degree of bargaining power downstream, ζ. Lemma 1 shows that under Assumption 1 and Assumption 2, wholesale prices and downstream prices are decreasing on the retailer's bargaining power ζ.

Lemma 1. *Wholesale prices and retail prices are decreasing in the bargaining power of the retailer. For $\lambda \geq 1$ and if Assumption 1 holds, then $\partial t / \partial \zeta < 0$ and $\partial p / \partial \zeta < 0$. For $\lambda < 1$ and if Assumption 1 and 2 hold, then $\partial t / \partial \zeta < 0$ and $\partial p / \partial \zeta < 0$.*

Proof. See Section A.1 □

We now establish general results for how market outcomes vary with the extent of alignment between the retailer and consumers, λ. When $\lambda \geq 1$, the retailer sets its price to be equal to the wholesale price, $p = t$. Lemma 2 shows that in this case, the wholesale price and the retail price are independent of λ. When $\lambda < 1$, the wholesale price is not always decreasing with λ. The intuition is as follows: as λ goes up, the retailer would like to give away profits to increase output. In some cases, this allows the upstream firm to set a higher wholesale price. Regardless, Lemma 2 shows that retail prices are decreasing with λ under Assumptions 1, 2 and 3, which is the result of main interest in our context.

Lemma 2. *The retail price is weakly decreasing in the weight given to consumer surplus, λ. In particular, for $\lambda \geq 1$ we show that $\partial p/\partial\lambda = 0$ and $\partial t/\partial\lambda = 0$. For $\lambda < 1$ and if Assumptions 1, 2, and 3 hold, then $\frac{\partial p}{\partial \lambda} < 0$.*

Proof. See Section A.1. □

Parametric Examples

Lemmas 1 and 2 provide general conditions under which retail prices are lower when retailers have more bargaining power, and when retailers are more aligned with consumers. These conditions hold for multiple families of demand that satisfy combinations of Assumptions 1, 2, and 3. To provide examples for these results, Lemmas 3-7 show that retail prices are weakly decreasing with ζ and λ for commonly used families of demand functions.

Lemma 3 (CES demand)**.** *Consider the CES demand function of the form $q = p^{\alpha}$, with $\alpha < -1$. With CES demand, wholesale and retail prices are weakly decreasing in the bargaining power downstream and in the weight given to consumer surplus. For $\lambda < 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} < 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$. For $\lambda \geq 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} = 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$.*

Proof. See Section A.1. □

Lemma 4 (Constant marginal revenue)**.** *Consider a demand function that features a constant marginal revenue curve $q = \frac{1}{p-a}$ (CMR demand). With CMR demand, wholesale prices and retail prices are weakly decreasing in the bargaining power downstream and in the weight given to consumer surplus. For $\lambda < 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} < 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} < 0$. For $\lambda \geq 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} = 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$.*

Proof. See Section A.1. □

Lemma 5 (Logit demand)**.** *Consider a logit demand function $q = \frac{e^{-\beta p}}{1+e^{-\beta p}}$. With logit demand, retailer prices are weakly decreasing in the retailer's bargaining power and in the weight given to consumer surplus. For $\lambda < 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} < 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$. For $\lambda \geq 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} = 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$.*

Proof. See Section A.1. □

Lemma 6 (Exponential demand)**.** *Consider an exponential demand function $q = e^{-\beta p}$. With exponential demand, retail prices are weakly decreasing in the retailer's bargaining power and in the weight given to consumer surplus. For $\lambda < 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} < 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$. For $\lambda \geq 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} = 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$.*

Proof. See Section A.1. □

Lemma 7 (ρ-linear demand)**.** *Consider a ρ-linear demand function $q = (a - bp)^{1/\rho}$. With $\rho-$linear demand, retail prices are weakly decreasing in the retailer's bargaining power and in the weight given to consumer surplus. For $\lambda < 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} < 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$. For $\lambda \geq 1$, $\frac{\partial p}{\partial \zeta} < 0$ and $\frac{\partial p}{\partial \lambda} = 0$, and in addition $\frac{\partial t}{\partial \zeta} < 0$ and $\frac{\partial t}{\partial \lambda} = 0$.*

Proof. See Section A.1. □

Additional Lemmas and Proofs

Assumption 3 and log-concavity

Lemma 8. *If q is twice differentiable and log-concave, then Assumption 3 holds: $\frac{q^2}{-q'} - CS \geq 0$.*

Proof. Since q is differentiable, q' exists and is finite. $\frac{q^2}{-q'} = 0$ and $CS = 0$ if $q = 0$. As $\lim_{p\to+\infty} q = 0$, $\lim_{p\to\infty} \frac{q^2}{-q'} - CS = 0$. Taking the derivatives of $f(p) \coloneqq \frac{q^2}{-q'} - CS$, we get:

$$f'(p) = \frac{-2qq'^2 + q^2q''}{q'^2} + q = \frac{-qq'^2 + q^2q''}{q'^2} = q\frac{-q'^2 + qq''}{q'^2} < 0,$$

as q is log-concave. So $f(p)$ is decreasing in p. From $\lim_{p\to\infty} f(p) = 0$ we get $f(p) \geq 0$. □

Decreasing Marginal Revenue

We provide an equivalent expression of decreasing marginal revenue for a twice-differentiable function.

Lemma 9. *If q is twice differentiable, then $2q'^2 - qq'' \geq 0$ if and only if q has decreasing marginal revenue.*

Proof. Rewrite marginal revenue MR as a function of p by inverse function theorem:

$$MR(p) = p + \frac{q(p)}{q'(p)}.$$

Taking the derivative with respect to p yields:

$$MR'(p) = 1 + \frac{q'^2 - qq''}{q'^2} = \frac{2q'^2 - qq''}{q'^2}.$$

so that marginal revenue is increasing in p if and only if $2q'^2 - qq'' \geq 0$. Since q is decreasing in p, marginal revenue is decreasing in q if and only if $2q'^2 - qq'' \geq 0$. □

Proof of Lemma 1

Case 1: $\lambda < 1$ In this case, the first order condition for the retailer holds and therefore:

$$F_2 \coloneqq p - t + (1-\lambda)\frac{q}{q'} = 0.$$

Taking the derivatives with respect to t yields:

$$\frac{dp}{dt} = -\frac{\frac{\partial F_2}{\partial t}}{\frac{\partial F_2}{\partial p}} = \frac{1}{\lambda + (1-\lambda)\frac{2(q')^2 - qq''}{(q')^2}}.$$

By Assumption 1, $\frac{dp}{dt} > 0$. Imposing condition F_2 on Equation (A.1), we obtain:

$$F_1 \coloneqq -\frac{q}{\frac{q'p'}{q} + \frac{1}{t-c}} + \frac{1-\zeta}{\zeta}\left[-(1-\lambda)\frac{q^2}{q'} + \lambda\,\mathrm{CS}\right] = 0,$$

such that:

$$\begin{aligned}
\frac{\partial F_1}{\partial \zeta} &= -\frac{1}{\zeta^2}\left[-(1-\lambda)\frac{q^2}{q'} + \lambda CS\right] \\
\frac{\partial F_1}{\partial t} &= -\frac{\frac{(2q'^2 - qq'')p'^2 - qq'p''}{q} + \frac{q'p'(t-c)+q}{(t-c)^2}}{\left(\frac{q'p'}{q} + \frac{1}{t-c}\right)^2} + \frac{1-\zeta}{\zeta}\left[-(1-\lambda)\frac{qp'(2q'^2 - qq'')}{q'^2} - \lambda qp'\right].
\end{aligned}$$

It follows immediately that $\frac{\partial F_1}{\partial \zeta} < 0$. The sign of p'' is determined by $\frac{d\frac{qq''}{q'^2}}{dp}$ since:

$$p'' = \frac{(1-\lambda)\frac{d\frac{qq''}{q'^2}}{dp}\frac{dp}{d\lambda}}{[(2-\lambda) - (1-\lambda)\frac{qq''}{q'^2}]^2}.$$

From Assumption 2, $\frac{d\frac{qq''}{q'^2}}{dp} \geq 0$, and therefore $p'' \geq 0$. The first term of $\frac{\partial F_1}{\partial \zeta}$ is weakly negative, and the second term is negative, so $\frac{\partial F_1}{\partial \zeta} < 0$. Therefore $\frac{\partial t}{\partial \zeta} = -\frac{\frac{\partial F_1}{\partial \zeta}}{\frac{\partial F_1}{\partial t}} < 0$. In addition, we get $\frac{\partial p}{\partial \zeta} < 0$ since $\frac{dp}{dt} > 0$.

Case 2: $\lambda \geq 1$ With a sufficiently high weight given to consumer surplus, in particular when $\lambda > 1$, the retailer will set the price equal to its marginal cost, as shown by equation (A.2).

The Nash bargaining first-order condition in equation (A.1) becomes:

$$F := \frac{-q^2(t)(t-c)}{(t-c)q'(t)+q(t)} + \frac{1-\zeta}{\zeta}CS(t) = 0.$$

Taking the partial derivative with respect to ζ yields:

$$\frac{\partial F}{\partial \zeta} = -\frac{1}{\zeta^2}CS(t) < 0; \quad \frac{\partial F}{\partial t} = -\frac{\frac{2q'^2-q''q}{q} + \frac{q'(t-c)+q}{(t-c)^2}}{\left(\frac{q'}{q} + \frac{1}{t-c}\right)^2} - \frac{1-\zeta}{\zeta}q < 0.$$

and it follows that under Assumption 1 that $\partial t/\partial \zeta < 0$.

Proof of Lemma 2

Case 1: $\lambda \geq 1$ For $\lambda \geq 1$, the retailer sets price equal to marginal cost, $p = t$. The Nash bargaining first order condition is:

$$F := \frac{-q^2(t)(t-c)}{(t-c)q'(t)+q(t)} + \frac{1-\zeta}{\zeta}CS(t) = 0,$$

which does not contain λ, so that t does not depend on λ. Thus $\frac{\partial t}{\partial \lambda} = 0$. From $p = t$, we have $\frac{\partial p}{\partial \lambda} = 0$.

Case 2: $\lambda < 1; \zeta = 0$ In the special case in which $\zeta = 0$, the upstream firm acts as a monopoly and sets the wholesale price to maximize its profits. In this case, the upstream firm and retailer profit functions become:

$$\begin{aligned} \Pi_U &= (t-c)q(p) \\ V_D &= (p-t)q(p) + \lambda\,\mathrm{CS}(p), \end{aligned}$$

and the upstream firm and retailer first order conditions become:

$$\begin{aligned} F_1 &:= (t-c)q'p' + q = 0 \\ F_2 &:= p - t + (1-\lambda)\frac{q}{q'} = 0, \end{aligned}$$

such that from the retailer's first order condition we obtain:

$$p' = -\frac{\frac{\partial F_2}{\partial t}}{\frac{\partial F_2}{\partial p}} = \frac{1}{1+(1-\lambda)\frac{q'^2-qq''}{q'^2}},$$

which we plug into the upstream firm's first order condition to rewrite F_1 as:

$$q + (t-c)q' \frac{1}{1 + (1-\lambda)\frac{q'^2 - qq''}{q'^2}} = 0.$$

By combining the two first order conditions, we get:

$$F := \begin{bmatrix} q + (t-c)q' \frac{1}{1+(1-\lambda)\frac{q'^2-qq''}{q'^2}} \\ p - t + (1-\lambda)\frac{q}{q'} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Note that:

$$\begin{bmatrix} \frac{\partial F_1}{\partial \lambda} \\ \frac{\partial F_2}{\partial \lambda} \end{bmatrix} = \begin{bmatrix} (t-c)q'p'^2 \frac{q'^2-qq''}{q'^2} \\ -\frac{q}{q'} \end{bmatrix},$$

and the Jacobian matrix of F is:

$$J := \begin{bmatrix} \frac{\partial F_1}{\partial t} & \frac{\partial F_1}{\partial p} \\ \frac{\partial F_2}{\partial t} & \frac{\partial F_2}{\partial p} \end{bmatrix} \begin{bmatrix} q'p' & q' + (t-c)\left[q''p' + q'p'^2(1-\lambda)\frac{d\frac{qq''}{q'^2}}{dp}\right] \\ -1 & \frac{1}{p'} \end{bmatrix},$$

while the determinant of J is:

$$\begin{aligned} det(J) &= q' + (t-c)\left[q''p' + q'p'^2(1-\lambda)\frac{d\frac{qq''}{q'^2}}{dp}\right] + \frac{q'p'}{p'} \\ &= \frac{2q'^2 - qq''}{q'} - (1-\lambda)qp'\frac{d\frac{qq''}{q'^2}}{dp}. \end{aligned}$$

From Assumption 1 and Lemma 9, $2q'^2 - qq'' \geq 0$. This yields $\frac{2q'^2-qq''}{q'} \leq 0$. From assumption 2,$\frac{d\frac{qq''}{q'^2}}{dp} \geq 0$. So $det(J) \leq 0$.

The inverse matrix of J is:

$$J^{-1} = \frac{1}{det(J)} \begin{bmatrix} \frac{1}{p'} & -q' - (t-c)\left[q''p' + q'p'^2(1-\lambda)\frac{d\frac{qq''}{q'^2}}{dp}\right] \\ 1 & q'p' \end{bmatrix},$$

and using the implicit function theorem we show that:

$$\begin{aligned} \frac{\partial p}{\partial \lambda} &= -\frac{1}{det(J)}\left[q'p' \cdot \left(-\frac{q}{q'}\right) + (t-c)q'p'^2 \frac{q'^2 - qq''}{q'^2}\right] \\ &= \frac{1}{det(J)} qp' \frac{2q'^2 - qq''}{q'^2} < 0. \end{aligned}$$

Case 3: $\lambda < 1; \zeta < 1$ Rewrite the first order condition for the bargaining problem as:

$$\frac{\frac{\partial \pi_u}{\partial t}}{\pi_u} - \frac{\zeta}{1-\zeta}\frac{q}{V_D} = 0 \Longrightarrow (1-\zeta)F_2 - \zeta q \frac{\pi_u}{V_D} = 0.$$

where $F_2 := q + (t-c)q'(p)p'(t)$. The first order condition of the retailer is:

$$F_1 := p - t + (1-\lambda)\frac{q}{q'} = 0.$$

Combining both conditions yields:

$$F := \begin{bmatrix} F_1 \\ (1-\zeta)F_2 - \zeta q \frac{\pi_u}{V_D} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix},$$

for which the partial derivative with respect to λ is:

$$\frac{\partial F}{\partial \lambda} = \begin{bmatrix} \frac{\partial F_1}{\partial \lambda} \\ (1-\zeta)\frac{\partial F_2}{\partial \lambda} - \zeta q \frac{\partial \frac{\pi_u}{V_D}}{\partial \lambda} \end{bmatrix}.$$

and the Jacobian is:

$$J = \begin{bmatrix} \frac{\partial F_1}{\partial t} & \frac{\partial F_1}{\partial p} \\ (1-\zeta)\frac{\partial F_2}{\partial t} - \zeta q \frac{\partial \frac{\pi_u}{V_D}}{\partial t} & (1-\zeta)\frac{\partial F_2}{\partial p} - \zeta q' \frac{\pi_u}{V_D} - \zeta q \frac{\partial \frac{\pi_u}{V_D}}{\partial p} \end{bmatrix},$$

for which the determinant is:

$$\det(J) = (1-\zeta)\left(\frac{\partial F_1}{\partial t}\frac{\partial F_2}{\partial p} - \frac{\partial F_1}{\partial p}\frac{\partial F_2}{\partial t}\right) + \zeta \left[q' \frac{\pi_u}{V_D} + q \frac{\partial \frac{\pi_u}{V_D}}{\partial p} + q \frac{\partial \frac{\pi_u}{V_D}}{\partial t}\frac{1}{p'(t)} \right].$$

We know that when $\zeta = 0$, then $\frac{\partial F_1}{\partial t}\frac{\partial F_2}{\partial p} - \frac{\partial F_1}{\partial p}\frac{\partial F_2}{\partial t} > 0$. So we focus on $M := q'\frac{\pi_u}{V_D} + q\frac{\partial \frac{\pi_u}{V_D}}{\partial P} + q\frac{\partial \frac{\pi_u}{V_D}}{\partial t}\frac{1}{p'(t)}$, which can be simplified to $M = \frac{q}{\zeta V_D p'}\left[(1+\zeta)q'p'(t-c) + q\right]$. This yields:

$$\det(J) = (1-\zeta)\left(\frac{\partial F_1}{\partial t}\frac{\partial F_2}{\partial p} - \frac{\partial F_1}{\partial p}\frac{\partial F_2}{\partial t}\right) + \frac{q}{V_D p'}\left[(1+\zeta)q'p'(t-c) + q\right].$$

where the first term is greater than 0 given $\zeta = 0$. The second term is decreasing in ζ. When $\zeta \to 1$, $t \to c$ because the wholesaler's profit has zero weight in the bargaining stage. Thus the second term is equal to $\frac{q^2}{V_D p'} > 0$. So $|J| > 0$ for all ζ. Thus, the inverse of the Jacobian is:

$$J^{-1} = \frac{1}{det(J)}\begin{bmatrix} (1-\zeta)\frac{\partial F_2}{\partial p} - \zeta q' \frac{\pi_u}{V_D} - \zeta q \frac{\partial \frac{\pi_u}{V_D}}{\partial p} & -\frac{\partial F_1}{\partial p} \\ -(1-\zeta)\frac{\partial F_2}{\partial t} + \zeta q \frac{\partial \frac{\pi_u}{V_D}}{\partial t} & \frac{\partial F_1}{\partial t} \end{bmatrix}.$$

Using these results, we can write the partial derivative of retail price with respect to λ as:

$$\frac{\partial p}{\partial \lambda} = -\frac{1}{\det(J)}\left[(1-\zeta)\left(-\frac{\partial F_2}{\partial t}\frac{\partial F_1}{\partial \lambda} + \frac{\partial F_1 \partial F_2}{\partial t}\right) + \zeta q\left(\frac{\partial \frac{\pi_u}{V_D}}{\partial t}\frac{\partial F_1}{\partial \lambda} - \frac{\partial \frac{\pi u}{v_D}}{\partial \lambda}\frac{\partial F_1}{\partial t}\right)\right],$$

where since $\zeta = 0$ we know that $-\frac{\partial F_2}{\partial t}\frac{\partial F_1}{\partial \lambda} + \frac{\partial F_1 \partial F_2}{\partial t} > 0$, so we can focus on the sign of the last term:

$$\frac{\partial \frac{\pi_u}{V_D}}{\partial t}\frac{\partial F_1}{\partial \lambda} - \frac{\partial \frac{\pi u}{v_D}}{\partial \lambda}\frac{\partial F_1}{\partial t} = \left[-\frac{q}{q'}\frac{q\,(V_D + \pi_u)}{{V_D}^2} - \frac{\mathrm{CS}\,\pi_u}{{V_D}^2}\right] = \frac{1}{V_D^2}\left[-\frac{q^2}{q'}V_D + \pi_u(-\frac{q^2}{q'} - CS)\right],$$

and from Assumption 3, we have $N > 0$. Therefore, $\frac{\partial p}{\partial \lambda} < 0$.

<u>*Case 4:* $\lambda < 1; \zeta = 1$</u> In this case, the upstream firm will set the wholesale price equal to the marginal cost, $t = c$. Equation (A.2) can be written as:

$$F \coloneqq (p-c)q' + (1-\lambda)q = 0.$$

from where by taking partial derivatives with respect to p we get:

$$\begin{aligned}\frac{\partial F}{\partial p} &= [q' + (p-c)q''] + (1-\lambda)q' = \lambda q' + (1-\lambda)\frac{2q'^2 - qq''}{q'} < 0\\ \frac{\partial F}{\partial \lambda} &= -q < 0,\end{aligned}$$

such that under Assumption 1, $\frac{\partial p}{\partial \lambda} = -\frac{\frac{\partial F}{\partial \lambda}}{\frac{\partial F}{\partial t}} < 0$.

Proof for Lemma 3 (CES demand)

Notice that the CES function is not quasi-concave. Note also that Assumption 1 holds:

$$2q'^2 - qq'' = 2\alpha^2 p^{2\alpha-2} - \alpha(\alpha-1)p^{2\alpha-2} = \alpha(\alpha+1)p^{2\alpha-2} > 0,$$

and that Assumption 2 holds:

$$(\log q)' + (\log q'')' - 2(\log(-q'))' = \log(-\alpha) + \log(1-\alpha) - 2\log(-\alpha) = \log(1-\alpha) - log(-\alpha) > 0,$$

such that Lemma 1 implies that $\partial p/\partial \zeta < 0$ and $\partial t/\partial \zeta < 0$. However, Assumption 3 fails to hold since:

$$\frac{p^{2\alpha}}{-\alpha p^{\alpha-1}} + \frac{p^{\alpha+1}}{\alpha+1} = \frac{p^{\alpha+1}}{-\alpha} + \frac{p^{\alpha+1}}{\alpha+1} = \frac{-p^{\alpha+1}}{\alpha(\alpha+1)} < 0.$$

From equation (A.2) we get:

$$p = \frac{\alpha}{\alpha + (1-\lambda)}t,$$

and then from equation (A.1) we get:

$$(t-c)\alpha p^{\alpha-1}\frac{\alpha}{\alpha+(1-\lambda)}+p^{\alpha}=\frac{\zeta}{1-\zeta}\frac{t-c}{\frac{\lambda-1}{\alpha+1-\lambda}t+\lambda(-\frac{p}{\alpha+1})},$$

which can be simplified to:

$$\alpha^2(t-c)+\alpha t+\frac{\zeta}{1-\zeta}\left(\alpha^2+\alpha\right)(t-c)=0,$$

from where it follows that t is independent of λ, and so $\frac{\partial t}{\partial\lambda}=0$ and $\frac{\partial p}{\partial\lambda}=\frac{\alpha}{(\alpha+(1-\lambda))^2}t<0$.

Proof for Lemma 4 (CMR demand)

Notice that the CMR demand is not quasi-concave. Note also that Assumption 1 holds:

$$2q'^2-qq''=\frac{2}{(p-a)^4}-\frac{2}{(p-a)^4}=0,$$

and that Assumption 2 also holds given:

$$\log q+\log q''-2\log q'=-\log(p-a)+\log 2-3\log(p-a)+4\log(p-a)=\log 2$$

is constant on p, i.e., weakly increasing in p. Then from Lemma 1, $dp/d\zeta<0$ and $dt/d\zeta<0$. However, Assumption 3 fails to hold since:

$$\frac{q^2}{-q'}-CS=1+\log(p-a).$$

We now check the sign of $\frac{\partial p}{\partial\lambda}$ when $\lambda\leq 1$. The first order condition for the Nash problem in equation (A.1)) implies:

$$F\coloneqq\frac{\zeta}{1-\zeta}\left(\frac{t-c}{c-a}\right)-\frac{1-\lambda}{\lambda}-\log\lambda+\log(t-a)-C=0,$$

where C is an arbitrary constant that nonetheless determines the price. Taking partial derivatives yields:

$$\begin{aligned}\frac{\partial F}{\partial t}&=\frac{\zeta}{(1-\zeta)(c-a)}+\frac{1}{t-a}>0\\ \frac{\partial F}{\partial\lambda}&=-\frac{-\lambda+(1-\lambda)}{\lambda^2}-\frac{1}{\lambda}=\frac{1-\lambda}{\lambda^2}>0.\end{aligned}$$

Using the implicit function theorem:

$$\frac{\partial t}{\partial \lambda} = -\frac{\frac{\partial F}{\partial \lambda}}{\frac{\partial F}{\partial t}} < 0.$$

and plugging these terms back into p yields:

$$\frac{\partial p}{\partial \lambda} = \frac{1}{\lambda}\frac{\partial t}{\partial \lambda} - \frac{1}{\lambda^2}(t-a) < 0.$$

When $\lambda > 1$, $p = t$. t is not affected by λ. So $\frac{\partial p}{\partial \lambda} = \frac{\partial t}{\partial \lambda} = 0$.

Proof for Lemma 5 (Logit demand)

The logit demand is log-concave, since:

$$\begin{aligned} q'^2 - qq'' &= \beta^2 q^2(1-q)^2 - \beta^2 q^2(1-q)(1-2q) \\ &= \beta^2 q^2(1-q)(1-q-1+2q) = \beta^2 q^3(1-q) > 0. \end{aligned}$$

so that Assumptions 1 and 3 hold. In addition, Assumption 2 holds, since:

$$\frac{qq''}{q'^2} = \frac{\beta^2 q^2(1-q)(1-2q)}{\beta^2 q^2(1-q)^2} = \frac{1-2q}{1-q} = 1 - \frac{q}{1-q}$$

is decreasing in q, and thus increasing in p.

Proof for Lemma 6 (Exponential demand)

The exponential function is log-concave since:

$$q'^2 - qq'' = \beta^2 e^{-2\beta p} - \beta^2 e^{-\beta p} \cdot e^{-\beta p} = 0,$$

so that Assumptions 1 and 3 hold. In addition, Assumption 2 also holds, since:

$$(\log q)' + (\log q'')' - 2(\log(-q'))' = \beta + \beta - 2\beta = 0.$$

Proof for Lemma 7 (ρ-linear Demand)

The ρ-linear function is log-concave since:

$$q'^2 - qq'' = b^2 \frac{1}{\rho}(a-bp)^{2/\rho - 2} > 0,$$

so that Assumptions 1 and 3 hold. In addition, Assumption 2 also holds, since:

$$(\log q)' + (\log q'')' - 2(\log(-q'))' = \frac{-b}{a-bp}\left(\frac{1}{\rho} + \frac{1}{\rho} - 2 - 2\left(\frac{1}{\rho} - 1\right)\right) = 0.$$

A.2 Experimental evidence on shopping behavior

Our experiment provided consumers with information on the availability of public pharmacies as an affordable alternative for purchasing drugs. This appendix studies whether consumers learned about the availability and attributes of public pharmacies, and whether knowing about them changed their shopping behavior in the short term. We estimate the equation:

$$y_i = \beta T_i + X_i'\gamma + \eta_{c(i)} + \varepsilon_i \tag{A.3}$$

where y_i is the outcome of interest; T_i indicates whether a consumer was treated; X_i is a vector of controls that includes the dependent variable at baseline along with consumer age, education, gender, and indicators for whether the consumer is covered by public insurance and whether a household member suffers a chronic condition; $\eta_{c(i)}$ are county fixed effects. The coefficient β measures the average treatment effect of our informational intervention.

Information about public pharmacies rendered consumers more aware of their availability and attributes. Panel A in Table A.4.9 displays these results. Columns (1) and (2) show that information increased awareness about the availability of the public pharmacy by 7 percentage points, from a baseline level of 77 percent. Moreover, columns (4) and (5) show that information shifted consumer perceptions about drug prices at public pharmacies, which is their most salient attribute. In particular, perceived public pharmacy prices decreased by 9 percent as a result of the intervention. We also find that perceived waiting time for receiving drugs at the public pharmacy increased, which is their main disadvantage relative to private pharmacies. In particular, perceived waiting time increased by 20 percent.[2] These results are consistent with consumers becoming aware of public pharmacies and their competitive advantages and disadvantages relative to private pharmacies as public pharmacies enter local markets.

Consumers also seem to have reacted to the intervention in terms of shopping behavior. Panel B in Table A.4.9 displays results from linear probability models for enrollment in the public pharmacy, the decision to purchase, and the plan to use the pharmacy in the future. Although estimates are imprecise, they are positive and economically meaningful. The point estimate in column (2) indicates a 2-percentage-points increase in enrollment with public pharmacies by treated households—almost a 30 percent increase relative to the mean of the

[2] We address concerns related to sample attrition by reporting bounds suggested by Lee (2009) in Table A.4.9-A. In all cases, point estimates for both the lower and upper bound have the same sign as our estimated treatment effects. However, in some cases, the point estimate of the bound is not statistically different from zero, which implies that under relatively negative attrition scenarios, our treatment effects are not distinguishable from zero.

control group. The results in column (5) imply a 2.3-percentage-points increase in purchases in public pharmacies by treated households—more than an 80 percent increase relative to a baseline share of 2.8 percent in the control group. Finally, column (8) shows that our intervention increased the extent to which households plan to use the public pharmacy by 5 percentage points, which is as much as 10 percent relative to the baseline level for the control group.

Households with members who suffer chronic conditions react more strongly to the treatment. Columns (3), (6), and (9) study heterogeneity along this margin. All effects are larger for households with chronic conditions, although the differences are not statistically significant. Moreover, the treatment effects on effective and planned purchases are marginally statistically significant for consumers with chronic conditions. Consumers with chronic conditions are more likely to periodically shop for drugs and thus the group for which short-term effects are more likely to be detectable. Moreover, in many cases, public pharmacies prioritize the provision of drugs to treat chronic conditions, and thus the information in our intervention may be less relevant for consumers without any household member with a chronic condition. Treatment effects on consumers without a household member with a chronic condition are indeed close to zero across outcomes.[3]

These results suggest that as public pharmacies enter local markets, consumers become aware of their entry, their relative advantages in terms of lower prices, and their relative disadvantages in terms of convenience. Moreover, our findings suggest that consumers value the availability of public pharmacies and some—particularly those affected by a chronic condition—substitute toward public pharmacies to take advantage of their lower drug prices.

A.3 The price effects of competition by public pharmacies

In this section, we develop a simple model of consumer choice and firm competition based on Chen and Riordan (2008). The goal is to illustrate the conditions under which the entry of an additional firm to a market induces an increase or a decrease in the prices set by an incumbent firm. The environment is simple but captures several features of our setting.

Setup

Environment. There is a population of consumers of size one that faces the discrete choice problem of purchasing from the incumbent, purchasing from the entrant, or not purchasing at all, which is the outside option. We denote these options by $j \in \{I, E, O\}$, respectively. After normalizing the value of the outside option to 0, the value that consumer i gets from

[3] We report Lee bounds in Panel B in Table A.4.9 to address concerns about attrition. We find that point estimates for both the lower and upper bound for all outcomes have the same sign as our estimated treatment effects, although some of those bounds are not statistically different from zero.

each option is

$$\begin{aligned} u_{iI} &= v_{iI} - p_r \\ u_{iE} &= v_{iE} - p_u \\ u_{iO} &= 0 \end{aligned}$$

where v_{ij} is the willingness to pay and p_j is the price of each option. Willingness to pay v_i is drawn from a differentiable joint distribution $H(v)$, and may feature average differences across firms, may be heterogeneous across consumers within each firm and may be correlated across firms. Consumers choose the option that gives the highest utility, so that the probability that consumer i chooses option j is

$$\sigma_{ij} = P(u_{ij} \geq u_{ik} \quad \forall k)$$

which induces demand functions

$$s_j = \int \sigma_{ij} h(v) dv$$

which naturally depend on the set of firms in the market.

On the supply side, the incumbent firm I chooses p^I to maximize profits $s_I(p_I - c_I)$, which leads to an optimal monopoly price p_I^m before entry and an optimal duopoly price p_I^d after entry. The entrant firm is meant to capture public pharmacies in our setting. As such, we assume it sets prices at marginal cost to satisfy a break-even condition, which is $p_E^d = c_E$.[4]

When does entry increase prices?

The net price effects of entry depend on the relative importance of two competing forces: (i) the extent of substitution away from the monopolist, which imposes downward pressure on the incumbent price, and (ii) the extent to which the demand faced by the monopolist becomes steeper after entry, which imposes upward pressure on the incumbent price. To establish this intuition formally, we define $F(v_I)$ as the marginal distribution of willingness to pay for the incumbent and $G(v_E|v_I)$ as the distribution of willingness to pay for the entrant, conditional on that for the incumbent. Both of these distributions are defined under the joint distribution $H(v)$. With this notation, we can restate Theorem 1 in Chen and Riordan (2008), which establishes that—under a few fairly general assumptions—the incumbent price will increase upon entry if and only if

$$\int_{p_I^m}^{\infty} [G(v|v) - G(p_I^m|v)] f(v) dv \leq (p_I^m - c_I) \int_{p_I^m}^{\infty} [g(p_I^m|v) - g(v|v)] f(v) dv$$

and will otherwise decrease.

This condition compares the magnitude of the two effects of entry. The left-hand side of the equation is the *market share effect* of entry. This term measures the difference between

[4] All results hold for the case in which the entrant sets a profit-maximizing price.

the market share the incumbent gets from charging the monopoly price as a monopoly and as a duopoly; that is, before and after entry. The more market share the entrant takes away from the incumbent, the stronger the incentives the incumbent has to decrease price in response to entry. The right-hand side of the equation is the *price sensitivity effect* of entry. The magnitude of this effect depends on the difference between the slope of the residual demand curve the incumbent faces before and after entry. The steeper the demand curve after entry relative to before entry, the lower the extent of substitution away from the incumbent from marginal consumers upon entry, and therefore the stronger the incentive of the incumbent to increase price upon entry.

The relative strength of these effects will largely depend on the distribution of consumer preferences. For example, the likelihood of a price increase is higher with a negative correlation in willingness to pay. In this case, substitution toward the entrant is lower than under a distribution of preferences with a positive correlation. Moreover, those who substitute away from the incumbent are consumers with a relatively low willingness to pay for the incumbent among those who purchase from the incumbent before entry, which leads to a steeper residual demand curve after entry.

Simulation

In this section, we show the results of simulating the model. The goal is to show numerically how different parameter combinations yield different predictions regarding the sign of the price effect of entry.

Specification. A key input in the simulation is the joint distribution of willingness to pay for the firms in the market, H_v, which we assume follows a joint normal distribution:

$$\begin{pmatrix} v_I \\ v_E \end{pmatrix} \sim N\left(\begin{matrix} \delta_I \\ \delta_E \end{matrix} , \begin{matrix} \sigma_I^2 & \rho\sigma_I\sigma_E \\ \rho\sigma_I\sigma_E & \sigma_E^2 \end{matrix} \right)$$

where the mean willingness to pay for each firm is denoted by δ_I and δ_U. Differences between δ_I and δ_U capture vertical differentiation between firms and relative to the outside option. The dispersion of willingness to pay is captured by the variances σ_I^2 and σ_E^2, and the correlation between the willingness to pay for the incumbent and the entrant is captured by ρ. If the willingness to pay is positively correlated ($\rho > 0$), then consumers share similar preferences for both goods relative to the outside option. If instead willingness to pay is negatively correlated ($\rho < 0$), then consumers with a strong taste for one of the firms have a weak taste for the other firm. This parameter determines the extent to which the slope of demand the incumbent faces changes upon entry, which is key in determining the price effects of entry.

Simulation details. We simulate equilibrium prices and market shares for the environments before and after entry, for a range of parameters of the distribution of preferences. In particular, we set δ_I and δ_E so that $(\delta_I + \delta_E)/2 = 10$ and $\delta_I/\delta_E = k_\delta$ for a a grid of values for k_δ from 1 to 10; we set $\sigma_I = \sigma_E = \sigma$ and construct a grid of values for σ from 1 to 15; and we construct a grid of values for ρ between -1 and 1. We set marginal costs at $c_I = 6$ and c_E. For each combination of (k_δ, σ, ρ), we solve for optimal prices and resulting market shares before and after entry.

Results

Results on price effects and the distribution of preferences. Our simulations illustrate that consumer preferences over firms play a key role in determining the equilibrium effects of entry on prices. Figure A.4.5 displays results for simulations over a grid of values for heterogeneity in preferences σ and correlation in preferences across firms ρ, for relative mean preferences of $\delta_I/\delta_E = 4$.

These results show two main patterns. First, the price charged by the incumbent firm is more likely to increase when preferences for the incumbent are more negatively correlated with those for the entrant. A more negative correlation implies that marginal consumers who substitute toward the entrant are those with a low willingness to pay for the incumbent, which makes the residual demand curve of the incumbent steeper and therefore imposes incentives to increase prices. This is consistent with a stronger price-sensitivity effect. Second, the results show that the price charged by the incumbent is more likely to increase when there is more dispersion in preferences, which is partly driven by the fact that when such dispersion is low, the demand curve is flatter and there is limited scope for price increases.

In the context of our setting and empirical results, this simulation suggests that the correlation between preferences for private and public pharmacies is likely negative. This suggests that pharmacy attributes—beyond drug prices—play an important role in pharmacy choice. An attribute that could be important in generating this pattern is heterogeneity in consumer locations relative to pharmacies: Consumers who live closer to private pharmacies are likely to pay more for them than for public pharmacies, whereas the opposite may be true for consumers who live closer to public pharmacies.

Results on price effects and the relative quality of the entrant. In addition to studying the conditions under which incumbent prices increase upon entry, we use the model to illustrate the importance of vertical quality difference in determining the penetration of the entrant and the differences in prices between the incumbent and the entrant. Figure A.4.6 shows results from simulations of the model for a grid of values for the relative quality of the incumbent δ_I/δ_E, while keeping average quality across firms fixed. We fix the remainder of the distribution of preferences to values such that the price of the incumbent increases; namely, $\rho = -0.99$ and $\sigma = 2.55$.

We study the implications of vertical differentiation for market shares and prices. Panel A in Figure A.4.6 shows that while the entrant is able to steal market share from the incumbent, the extent of business stealing decreases substantially as the quality of the entrant relative to the incumbent decreases. Panel B in Figure A.4.6 shows that the incumbent price is higher when the quality of the entrant relative to the incumbent is lower. Furthermore, these results also show that the price effects of entry on the incumbent price depend on the relative quality of the entrant. The higher the relative quality of the entrant, the more likely the incumbent price will decrease upon entry.

These results are consistent with our descriptive evidence and main empirical findings. In Section 1.4, we documented that public pharmacies entered the market offering lower quality along several dimensions, which suggests that δ_I/δ_E is relatively large in our setting. These results indeed imply that entrants with low relative quality have low penetration, allow the incumbent to sustain higher prices, and make it more likely that the incumbent will increase prices.

A.4 Additional Figures and Tables

Figure A.4.1: Examples of private and public pharmacies

(a) Outside a private pharmacy

(b) Inside a private pharmacy

(c) Outside a public pharmacy

(d) Inside a public pharmacy

Notes: This figure displays photos of private and public pharmacies from the outside and inside. The private pharmacy in Panels (a) and (b) is a somewhat generic building and is one of the three leading chains in the market. The public pharmacy in Panels (c) and (d) is located in the city capital and is part of our experimental sample.

Figure A.4.2: Number of events per market and time dispersion

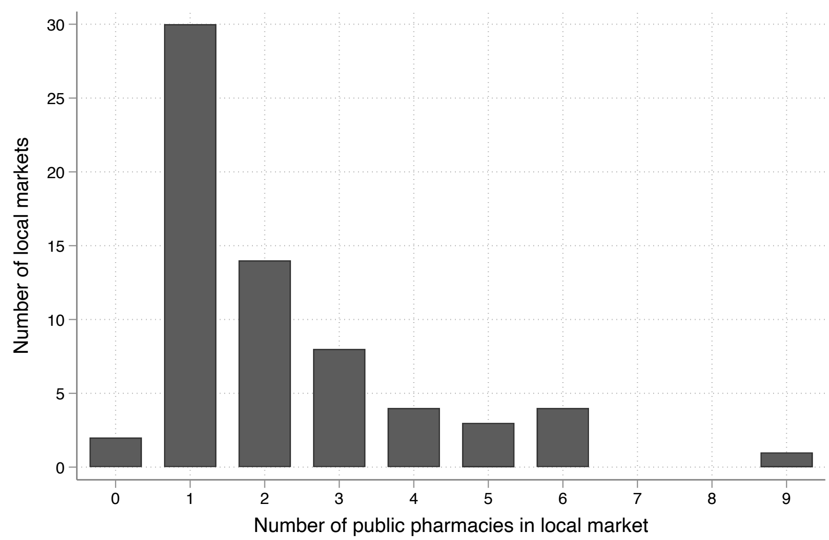

(a) Number of public pharmacy entries by market

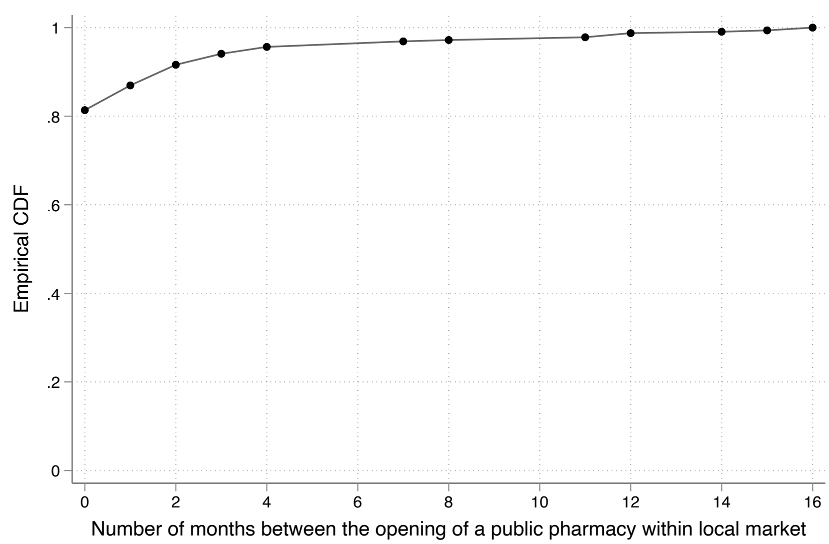

(b) Time between subsequent public pharmacy entries within a market

Notes: Panel (a) shows the distribution of the number of public pharmacies per local market. Panel (a) shows the cumulative distribution function of the dispersion of events within local markets. For example, more than 80 percent of events within the market occurred within the same month, which is by definition the case for markets with only one event.

Figure A.4.3: Impact of public pharmacies using only the first entry in a local market

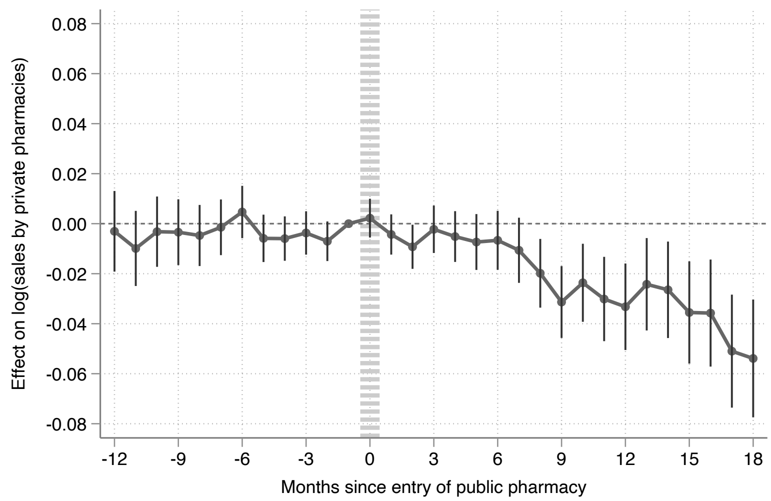

(a) Sales

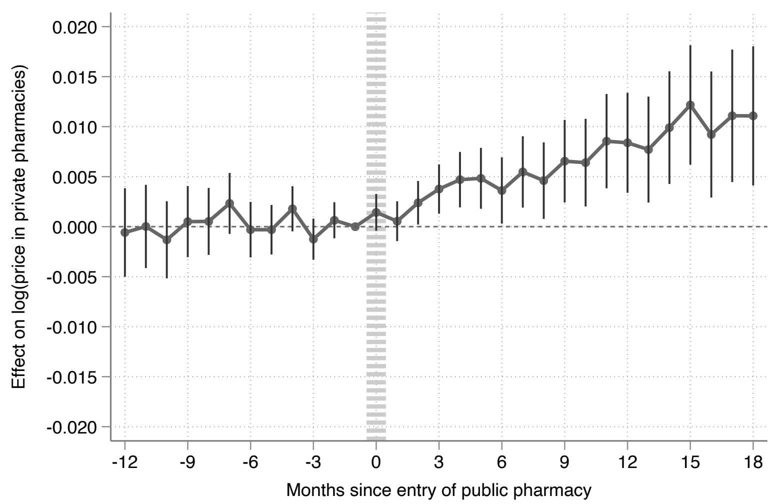

(b) Prices

Notes: These figures present event-study estimates of the impact of public pharmacies on private pharmacy sales in Panel (a), and on private pharmacy prices in Panel (b). The unit of observation is a molecule per market in a given month. The empirical strategy uses panel data for the period between 2014 and 2018 and exploits the staggered entry of public pharmacies from October 2015 onward in an event-study design. Treatment is defined as introduction of the *first* public pharmacy in the market. In Panel (a) the dependent variable is logged sales and in Panel (b) the dependent variable is logged prices. The x-axis indicates the month with respect to the opening of the public pharmacy, i.e., 18 means 18 months after the opening, and −12 means twelve months before the opening. Dots indicate estimated coefficients, and vertical lines indicate the corresponding 95 percent confidence intervals.

Figure A.4.4: Impact of public pharmacies in markets with events within less that 1 month

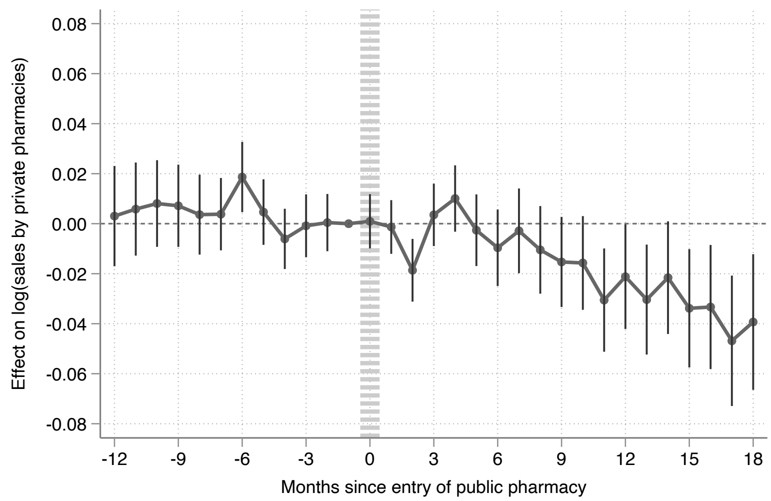

(a) Sales

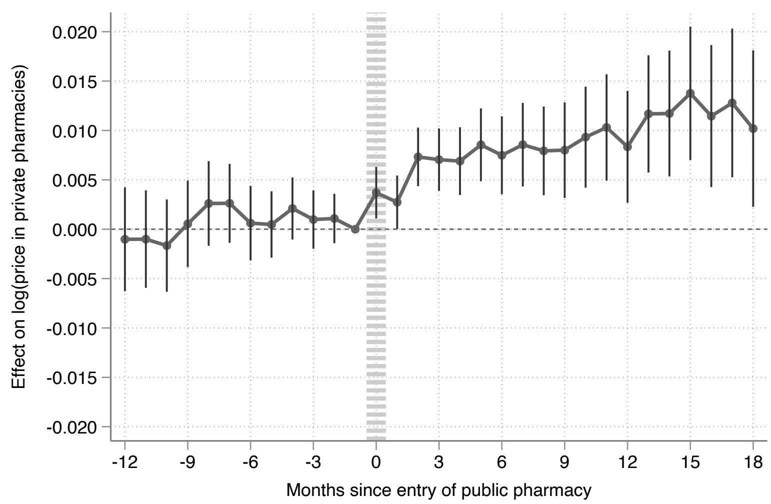

(b) Prices

Notes: These figures present event-study estimates of the impact of public pharmacies on private pharmacy sales in Panel (a), and on private pharmacy prices in Panel (b). The unit of observation is a molecule per market in a given month. The empirical strategy uses panel data for the period between 2014 and 2018 and exploits the staggered entry of public pharmacies from October 2015 onward in an event-study design. Treatment is defined as introduction of the *first* public pharmacy in the market. The sample only includes never-treated markets and markets with either one event or in which events are no more than 1 month apart. In Panel (a) the dependent variable is logged sales and in Panel (b) the dependent variable is logged prices. The x-axis indicates the month with respect to the opening of the public pharmacy, i.e., 18 means 18 months after the opening, and -12 means twelve months before the opening. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95 percent confidence intervals.

Figure A.4.5: Simulations for the price effects of entry

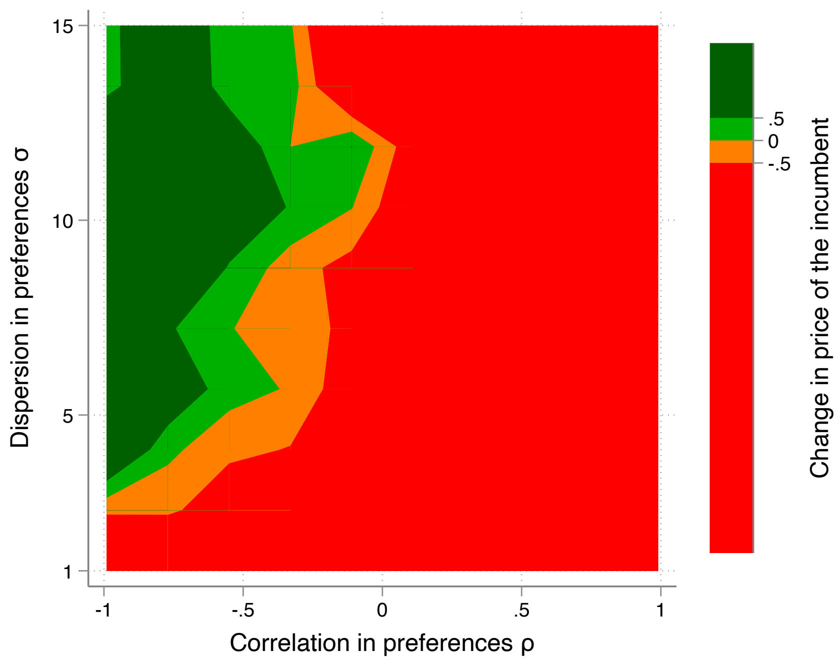

Notes: This figure plots simulated effects of entry on the price the incumbent charges, as discussed in A.3. The plot provides results for a grid of values of σ and ρ, under mean preferences for the incumbent and entrant $\delta_I/\delta_E = 4$, although the results are qualitatively similar for different values of the latter. The red region indicates that the incumbent price *decreases*, whereas the green region indicates that the incumbent price *increases* for each distributions of preferences, respectively.

Figure A.4.6: Simulations for the role of relative quality in equilibrium outcomes

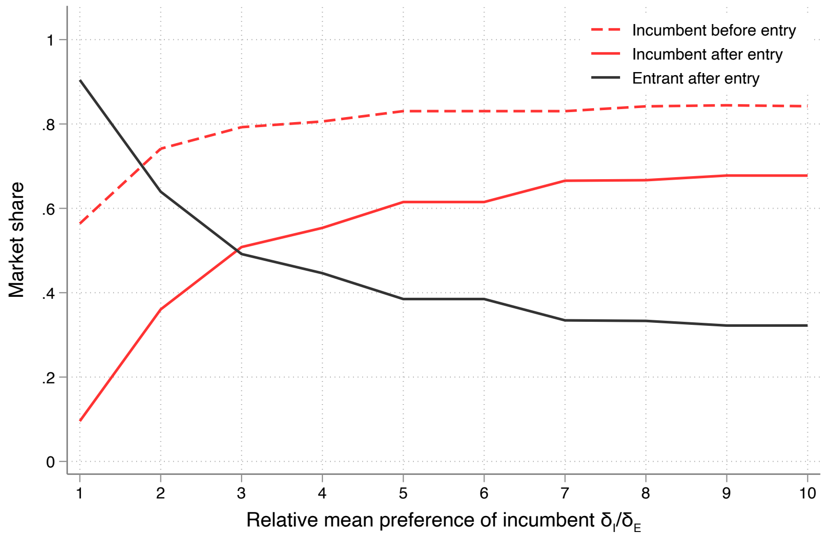

(a) Equilibrium market shares

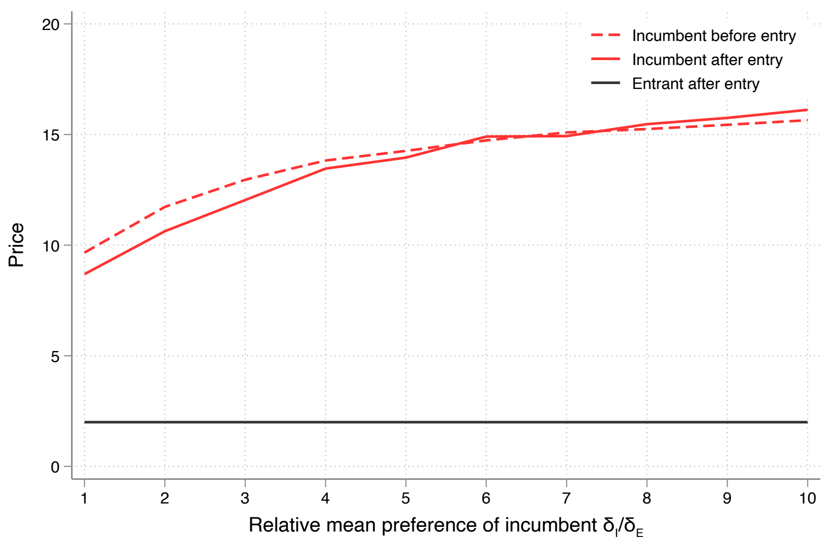

(b) Equilibrium prices

Notes: Both panels display equilibrium outcomes for the incumbent and entrant, before and after entry for a range of values for relative quality of the incumbent δ_I/δ_E, while keeping the average quality of both firms fixed. Panel (a) displays equilibrium market shares, whereas Panel (b) displays equilibrium prices. Incumbent outcomes are plotted in red, while entrant outcomes are plotted in black. Outcomes before entry are plotted in dashed lines, while outcomes after entry are plotted in dashed lines.

Figure A.4.7: Event study estimates for effects on municipal finance

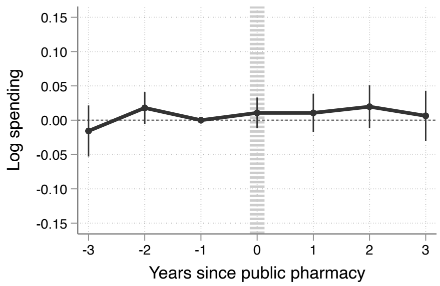

(a) Total spending

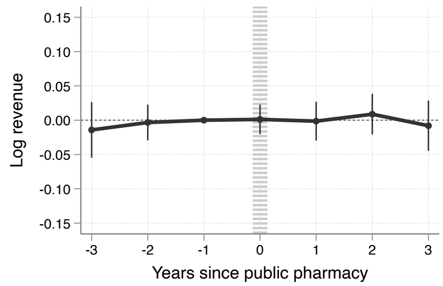

(b) Total revenue

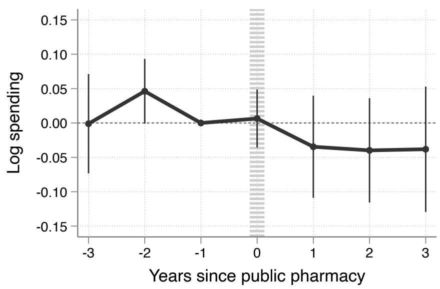

(c) Non-health spending

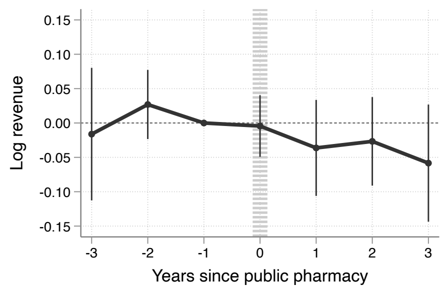

(d) Non-health revenue

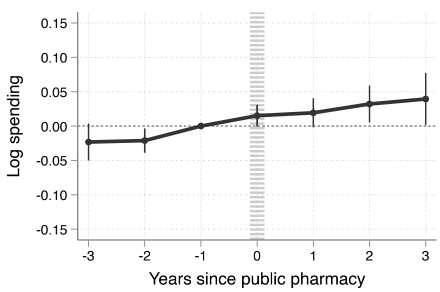

(e) Health spending

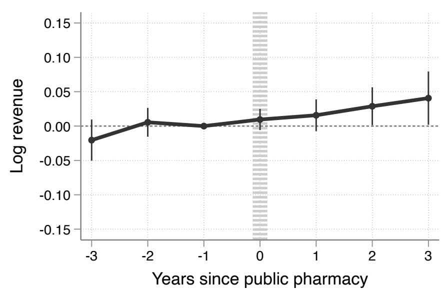

(f) Health revenue

Notes: These figures present event study estimates for the impact of public pharmacies on municipal finance using panel data for 2013-2019. Municipal spending and revenue are measured in monetary units per capita. Each plot displays results from an event study version of equation (1.3) given by:

$$y_{ct} = \sum_{k=-3}^{3} \delta_k D_{ct}^k + \theta_c + \lambda_t + \varepsilon_{ct},$$

where the outcomes are the measures of municipal finance (revenue, spending) and treatment dummies are defined with respect to the first year with a public pharmacy. All regressions include county fixed effects θ_c and year fixed effects λ_t. Dots indicate point estimates and vertical lines indicate the corresponding 95 percent confidence intervals.

Figure A.4.8: Event study estimates for effects on avoidable hospitalizations

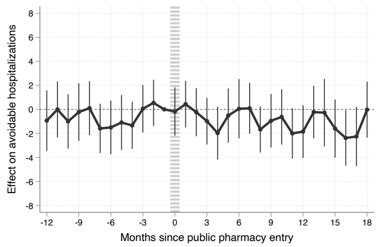

(a) Number of hospitalizations, all insurance

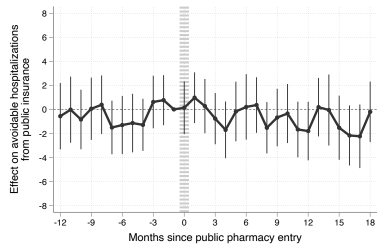

(b) Number of hospitalizations, public insurance

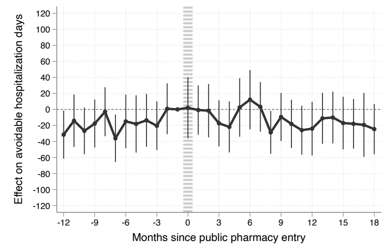

(c) Days of hospitalizations, all insurance

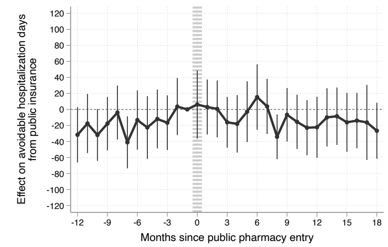

(d) Days of hospitalizations, public insurance

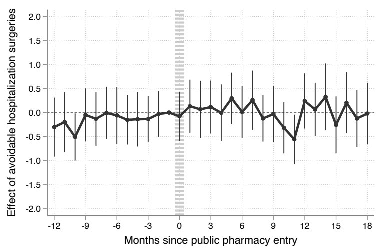

(e) Number of surgeries, all insurance

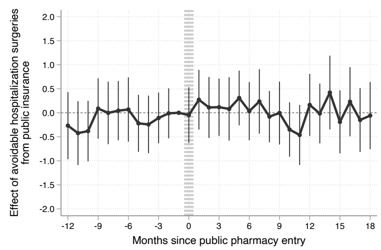

(f) Number of surgeries, public insurance

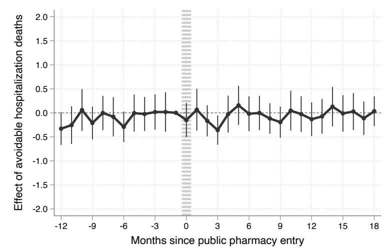

(g) Number of deaths, all insurance

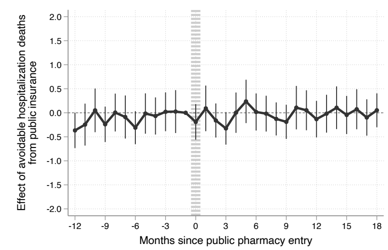

(h) Number of deaths, public insurance

Notes: Each plot displays results from an event study version of equation (1.3) given by:

$$y_{ct} = \sum_{k=-12}^{18} \delta_k D_{ct}^k + \theta_c + \lambda_t + \varepsilon_{ct},$$

where the outcomes are the same measures of avoidable hospitalization events as in Table 1.4 and treatment dummies D_{ct}^k are defined as a month t which is exactly k months after event time in county c. We normalize $\delta_{k=-1} = 0$, so we interpret all coefficients δ_k as the effect of a public pharmacy's opening on the dependent variable exactly k months after its entry. Dots indicate point estimates and vertical lines indicate the corresponding 95 percent confidence intervals.

Figure A.4.9: Other health outcomes, event study evidence

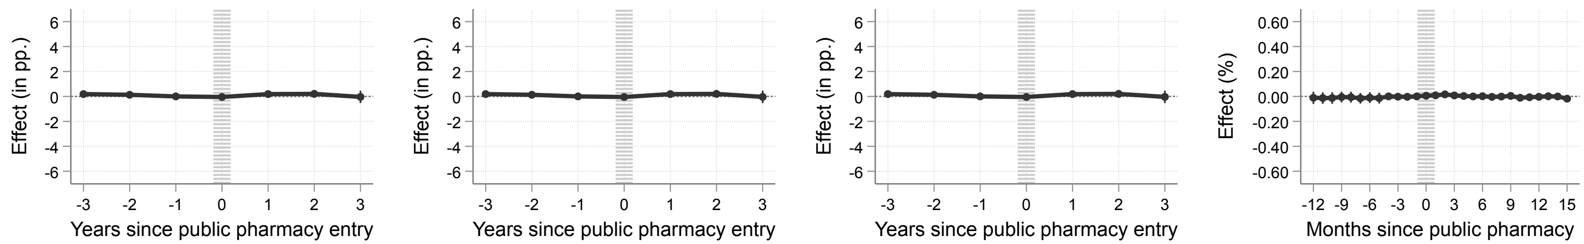

(a) Attendance, all schools (b) Attendance, public schools (c) Attendance, rural schools (d) Sick leaves, all

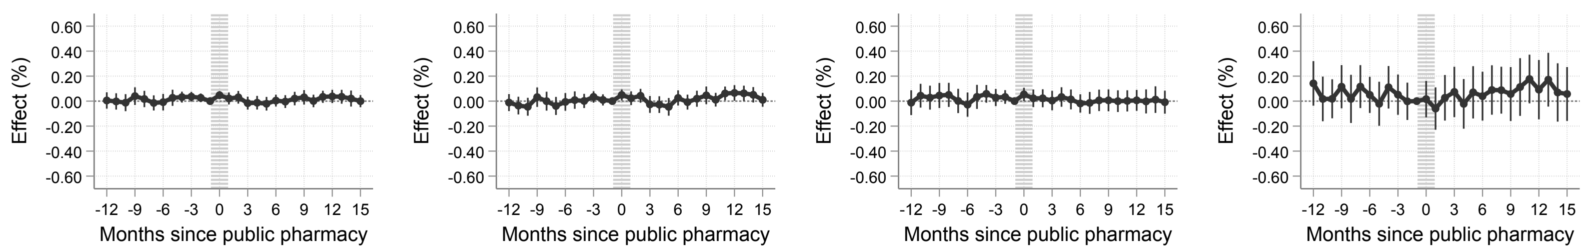

(e) Sick leaves, overall (f) Sick leaves, acute (g) Sick leaves, chronic (h) Sick leaves, diabetes

Notes: These figures present event study estimates for the impact of public pharmacies on school attendance using annual panel data for 2014-2019 and on sick leave using monthly panel data for 2015-2019. The former is administrative data from Ministry of Education and the latter is administrative data from the funding branch of the Ministry of Health (FONASA). Each plot displays results from an event study regression given by

$$y_{ct} = \sum_{k=\kappa_1}^{\kappa_2} \delta_k D_{ct}^k + \theta_c + \lambda_t + \varepsilon_{ct}$$

where the outcomes are school attendance in percentages ($\in [0, 100]$) and the number of sick leave per capita. Treatment indicators are defined with respect to the first year (panels a-c) or month (panels d-h) with a public pharmacy. All estimates include county fixed effects θ_c and year (or month-year) fixed effects λ_t. Each dot represents a coefficient and vertical lines indicate 95 percent confidence intervals.

Figure A.4.10: Informational treatment

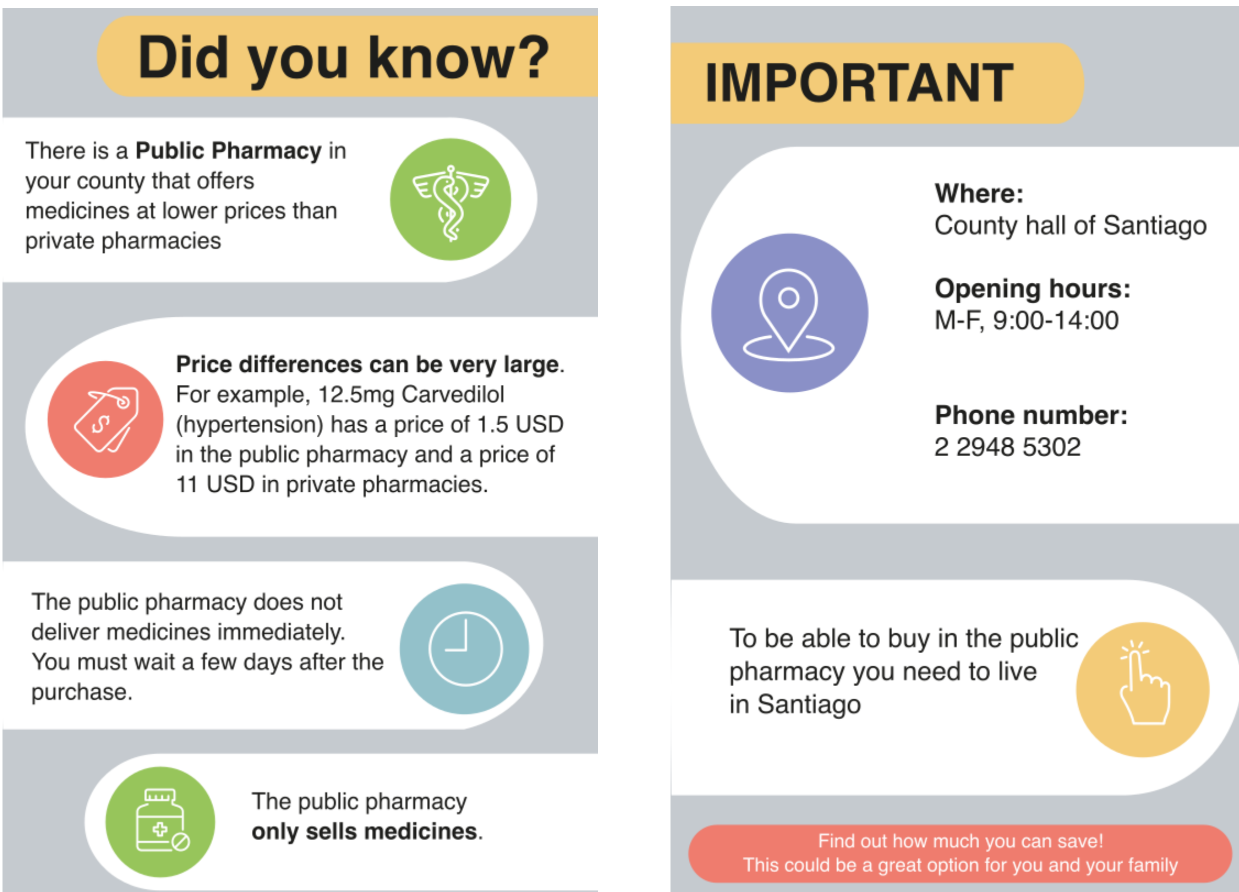

(a) Awareness and convenience (b) Search details

Notes: This figure displays the informational interventions delivered as part of the field experiment. Panel (a) displays the first part of the treatment, which aimed to increase awareness of the public pharmacy. It introduces the public pharmacy and states that it offers lower prices than private pharmacies and that it may take longer to deliver products. Panel (b) displays the second part, which aim to reduce search costs for participants by including detailed location and contact information for the public pharmacy, hours of operation, and eligibility requirements, tailored to the county of each participant.

Figure A.4.11: Location of pharmacies and consumers in experimental sample

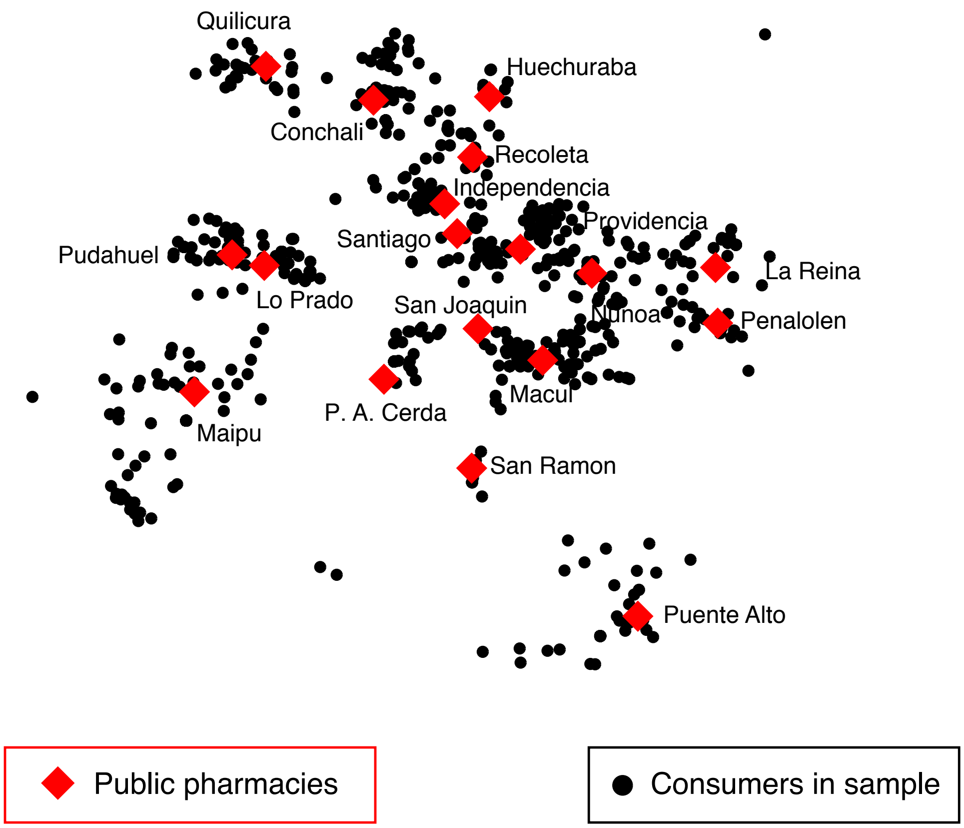

Notes: This figure displays the location of public pharmacies and consumers surveyed in the context of the field experiment. We surveyed 826 people at baseline outside randomly selected private pharmacies located in 18 counties within the city capital. All of these counties had a public pharmacy at the time of the baseline survey.

Figure A.4.12: Timeline of experiment events

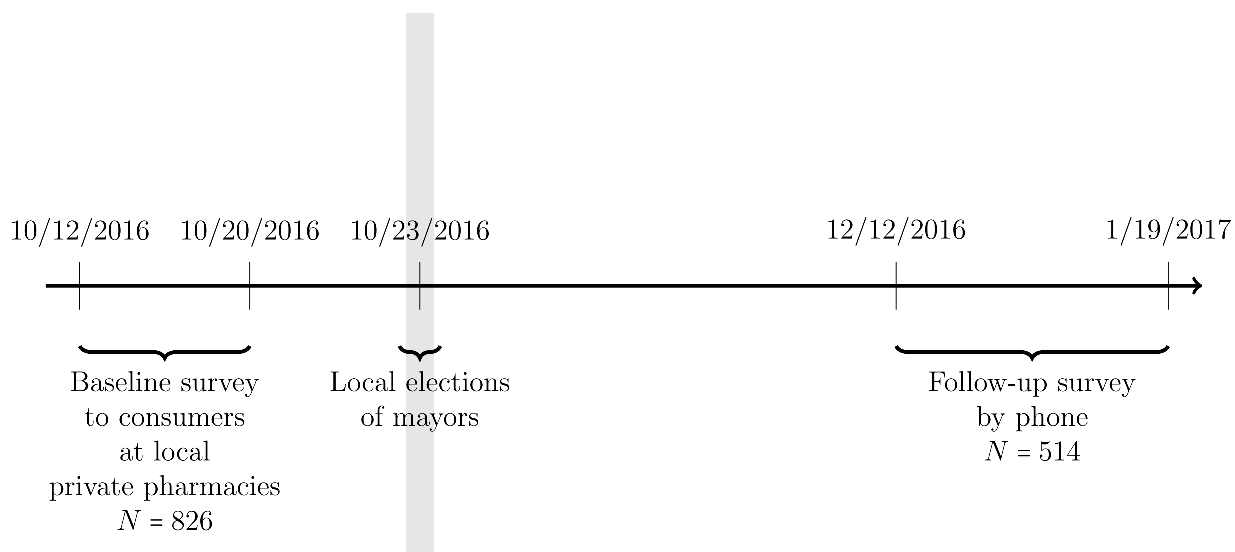

Notes: This timeline displays the main events in our field experiment. Baseline surveys were implemented outside randomly chosen private pharmacies in counties with a public pharmacy. Local elections are a single-day election held every 4 years in which citizens in all 344 counties vote for a mayor using simple majority rule. Follow-up surveys were implemented during a 1-month period to minimize attrition.

Table A.4.1: Within-county analysis of public pharmacy entry

	(1)	(2)	(3)	(4)
	1(Public pharmacy)			
Private pharmacies in 2014	0.022***	0.032***	0.030***	0.026***
	(0.004)	(0.006)	(0.005)	(0.004)
Schools in 2014	0.016***	0.011***	0.011***	0.007***
	(0.002)	(0.002)	(0.002)	(0.001)
Cell size is (in meters):	1,000	800	600	400
Cells	10,167	14,046	21,885	43,695
Mean of dependent variable	0.014	0.010	0.006	0.003
Mean of private pharmacies	0.198	0.142	0.088	0.049
County fixed effects	Yes	Yes	Yes	Yes

Notes: The unit of observation is a geographic cell within a county. We use all 146 counties with a public pharmacy operating by December 2018. Private pharmacies are measured in the year 2014, before the opening of public pharmacies. The estimating sample restricts attention to "populated cells," i.e., cells within the convex hull of schools in 2014. Different columns display results for different definitions of cell size, from 1,000×1,000 meters in column (1) to 400×400 meters in column (4). Standard errors clustered by county.

Table A.4.2: Effect on drug sales and prices in the private market

	(1) log(sales)	(2) log(price)
Panel A: Main estimates		
All sample (β^{jump})	0.0071*	0.0033***
	(0.0043)	(0.0011)
	[0.0074]	[0.0022]
All sample ($\beta^{\text{phase in}}$)	-0.0029***	0.0005***
	(0.0005)	(0.0001)
	[0.0008]	[0.0006]
R-squared	0.54	0.85
Panel B: Heterogeneity by chronic condition		
Molecules for chronic conditions ($\beta^{\text{jump}}_{\text{chronic}}$)	-0.0036	0.0029**
	(0.0050)	(0.0013)
	[0.0082]	[0.0023]
Molecules for non-chronic conditions ($\beta^{\text{jump}}_{\text{non-chronic}}$)	0.0219***	0.0037*
	(0.0079)	(0.0020)
	[0.0103]	[0.0027]
Molecules for chronic conditions ($\beta^{\text{phase in}}_{\text{chronic}}$)	-0.0028***	0.0005***
	(0.0005)	(0.0001)
	[0.0008]	[0.0006]
Molecules for non-chronic conditions ($\beta^{\text{phase in}}_{\text{non-chronic}}$)	-0.0030***	0.0005***
	(0.0007)	(0.0002)
	[0.0010]	[0.0007]
Panel C: Heterogeneity by relative public/private product variety		
High public-private variety ratio ($\beta^{\text{jump}}_{\text{high variety}}$)	-0.0110**	0.0037***
	(0.0048)	(0.0012)
	[0.0096]	[0.0025]
Low public-private variety ratio ($\beta^{\text{jump}}_{\text{low variety}}$)	0.0039	0.0031**
	(0.0057)	(0.0015)
	[0.0080]	[0.0023]
High public-private variety ratio ($\beta^{\text{phase in}}_{\text{high variety}}$)	-0.0035***	0.0007***
	(0.0005)	(0.0001)
	[0.0010]	[0.0006]
Low public-private variety ratio ($\beta^{\text{phase in}}_{\text{low variety}}$)	-0.0024***	0.0003**
	(0.0006)	(0.0002)
	[0.0009]	[0.0006]
Panel D: Heterogeneity by distance to private pharmacy		
Private pharmacies are close to public pharmacy ($\beta^{\text{jump}}_{\text{close}}$)	0.0010*	0.0044***
	(0.0053)	(0.0013)
	[0.0080]	[0.0025]
Private pharmacies are far from public pharmacy ($\beta^{\text{jump}}_{\text{far}}$)	0.0046	0.0023
	(0.0052)	(0.0014)
	[0.0108]	[0.0027]
Private pharmacies are close to public pharmacy ($\beta^{\text{phase in}}_{\text{close}}$)	-0.0034***	0.0003*
	(0.0006)	(0.0002)
	[0.0010]	[0.0007]
Private pharmacies are far from public pharmacy ($\beta^{\text{phase in}}_{\text{far}}$)	-0.0025***	0.0007***
	(0.0006)	(0.0002)
	[0.0009]	[0.0006]
Observations	681,120	649,885
Molecule-by-month FE	Yes	Yes
Molecule-by-market FE	Yes	Yes

Notes: This table presents our parametric estimates of the effects of public pharmacy entry on private market outcomes. We estimate the parameters β^{jump} and $\beta^{\text{phase in}}$ using an exposure difference-in-differences design that leverages the staggered introduction of public pharmacies in the panel data of molecules observed by market and month in the period 2014-2018. The parameter π^{jump} measures the immediate impact of public pharmacies and $\pi^{\text{phase in}}$ the additional impact by each year of operation. In Panel B, exposure to public pharmacies is interacted with an indicator for whether a molecule is targeted toward a chronic condition or not. In Panel C, exposure is interacted with an indicator for whether there is a high ratio of variety of products within a molecule in public pharmacies relative to private pharmacies. In Panel D, exposure is interacted with an indicator for whether in the local market private pharmacies were located relatively close to the public pharmacy. Standard errors clustered at the molecule-by-market level displayed in parentheses. *** p<0.01, ** p<0.05, * p<0.1. We also provide standard errors clustered at the local market level and are displayed in square brackets.

Table A.4.3: Municipal finance, full regression coefficients

	(1)	(2)	(3)	(4)	(5)	(6)
	Health services		Non-health services		All services	
	Spending	Revenue	Spending	Revenue	Spending	Revenue
π^{jump}: Public pharmacy	0.021*	0.002	0.010	-0.008	0.006	0.0003
	(0.011)	(0.010)	(0.036)	(0.039)	(0.017)	(0.016)
$\pi^{\text{phase in}}$: Public pharmacy × trend	0.008	0.010*	-0.023	-0.016	0.004	0.004
	(0.005)	(0.005)	(0.017)	(0.015)	(0.007)	(0.007)
Mean of dep. var. in 2014	170.36	167.09	525.32	563.07	695.68	730.15
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.968	0.964	0.939	0.937	0.977	0.976
Counties	321	321	322	322	321	322
Observations (county-years)	2,243	2,243	2,228	2,227	2,243	2,243

Notes: This table presents our estimates for the impact of public pharmacies on municipal finances. The health (columns 1-2) and non-health (columns 3-4) categories are mutually exclusive. Columns 5-6 correspond to "All services" provided by the county. We observe a panel of counties every year in the period 2013–2019 and exploit the staggered entry of pharmacies in a parametric event study analysis. The dependent variable is the logarithm of total spending (in U.S. dollars) per capita (2013 population) in odd columns and the logarithm of total revenue per capita in even columns. The parameter π^{jump} measures the immediate impact of public pharmacies and $\pi^{\text{phase in}}$ the additional impact by each year of operation. Standard errors clustered at the county level are displayed in parentheses. Significance level: *** p<0.01, ** p<0.05, * p<0.1.

Table A.4.4: Effect on avoidable hospitalizations associated with chronic diseases, full regression coefficients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Avoidable hospitalizations per 100,000 inhabitants							
	Number of hospitalizations		Days of hospitalizations		Number of surgeries		Number of deaths	
π^{jump}: Public pharmacy	0.332	0.586	13.897*	17.457**	0.187	0.235	-0.035	-0.015
	(0.489)	(0.532)	(7.091)	(7.838)	(0.149)	(0.166)	(0.069)	(0.077)
$\pi^{\text{phase in}}$: Public pharmacy × trend	-0.064	-0.085	-1.039*	-1.185*	-0.004	-0.008	0.007	0.007
	(0.049)	(0.052)	(0.578)	(0.646)	(0.013)	(0.014)	(0.006)	(0.006)
Health insurance	All	Public	All	Public	All	Public	All	Public
Mean of dep. var. in 2014	17.93	19.18	158.1	172.5	1.724	1.907	0.736	0.828
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.450	0.664	0.264	0.602	0.139	0.573	0.062	0.598
Counties	344	344	344	344	344	344	344	344
Observations (county-month-years)	28,320	28,320	28,320	28,320	28,320	28,320	28,320	28,320

Notes: This table presents our estimates for the impact of public pharmacies on avoidable health outcomes. The outcomes of interest are the number of hospitalizations (columns 1-2), days of hospitalizations (3-4), number of surgeries (columns 5-6), and number of deaths (columns 7-8). For each outcome, the first column uses the count of the outcome per 100,000 inhabitants in a county regardless of individual health insurance, and the second column restricts that count to individuals with publicly provided insurance (FONASA). We observe a panel of counties every month in the period 2013–2019 and exploit the staggered entry of pharmacies in a parametric event study analysis. The parameter π^{jump} measures the immediate impact of public pharmacies and $\pi^{\text{phase in}}$ the additional impact by each year of operation. We report the mean of the dependent variable for 2014 among counties that ever introduce a public pharmacy, the year before most public pharmacies entered the market. Standard errors clustered at the county level are displayed in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A.4.5: Public pharmacies and other health outcomes

	School attendance (county-year panel)			Sick leave (county-month panel)				
	All schools	Public schools	Rural schools	All diseases	Overall	Acute	Chronic	Diabetes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Public pharmacy 18-month effect	0.137	0.335	-0.081	0.002	-0.001	0.022	-0.011	-0.017
	(0.131)	(0.209)	(0.151)	(0.023)	(0.032)	(0.032)	(0.053)	(0.126)
Observations	2,070	2,031	1,802	20,002	18,194	17,441	15,096	8,509
Counties	345	345	301	346	341	336	336	311
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Avg. dependent variable	90.9	89.8	93.5	7.19	3.69	3.30	2.77	1.56

Notes: This table present difference-in-differences estimates for the impact of public pharmacies on school attendance using annual panel data for 2014-2019 and on sick leave using monthly panel data for 2015-2019. The former is administrative data from Ministry of Education and the latter is administrative data from the funding branch of the Ministry of Health (FONASA). Each column displays results from an event study regression given by:

$$y_{ct} = \theta_c + \lambda_t + \pi^{\text{jump}} PP_{ct} + \pi^{\text{phase in}} PP_{ct}(t - t_e^* + 1) + \varepsilon_{ct},$$

where the outcomes are school attendance in percentages ($\in [0, 100]$) and the number of sick leave per capita in county c and year t, PP_{ct} indicates the share of the year with a public pharmacy in county c, and $(t - t_e^*)$ measures the number of years since the opening of the public pharmacy. All regressions include county fixed effects θ_c and year (or month-year) fixed effects λ_t. Columns (1)-(3) use school absenteeism as dependent variable (years 2014-2019). Columns (4) and (5) use the logarithm of the total number of sick leave per 100,000 inhabitants (years 2015-2019). Standard errors are clustered at the county level. Significance level: *** p<0.01, ** p<0.05, * p<0.1.

Table A.4.6: Balance in covariates across attrition status

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Non-Attriters vs Attriters			Panel B: Non-Attriters		
Variable	Non-Attriters	Attriters	p-value $H_0: (1) = (2)$	Control	Treatment	p-value $H_0: (4) = (5)$
Monthly drug expenditure	75.44	78.48	0.57	78.05	73.56	0.54
	(71.93)	(70.37)		(75.50)	(69.31)	
Chronic condition in household	0.61	0.49	0.00	0.61	0.61	0.65
	(0.49)	(0.50)		(0.49)	(0.49)	
Age	46.70	44.60	0.09	46.62	46.77	0.62
	(16.67)	(18.08)		(16.84)	(16.57)	
Education higher than HS	0.53	0.52	0.89	0.54	0.52	0.72
	(0.50)	(0.50)		(0.50)	(0.50)	
Female	0.64	0.58	0.06	0.62	0.66	0.74
	(0.48)	(0.49)		(0.49)	(0.47)	
Public insurance	0.63	0.66	0.34	0.62	0.63	0.31
	(0.48)	(0.47)		(0.49)	(0.48)	
Day with internet (1-7)	5.26	5.43	0.40	5.12	5.35	0.37
	(2.84)	(2.71)		(2.92)	(2.78)	
Day with social media (1-7)	5.22	5.34	0.56	5.07	5.32	0.17
	(2.89)	(2.82)		(2.96)	(2.83)	
Employed	0.63	0.64	0.74	0.59	0.65	0.82
	(0.48)	(0.48)		(0.49)	(0.48)	
Supports incumbent	0.48	0.56	0.09	0.50	0.47	0.23
	(0.50)	(0.50)		(0.50)	(0.50)	
Voted in previous election	0.76	0.70	0.06	0.74	0.78	0.88
	(0.43)	(0.46)		(0.44)	(0.41)	
Knows public pharmacy	0.67	0.60	0.04	0.64	0.69	0.08
	(0.47)	(0.49)		(0.48)	(0.46)	
Perceived relative price of public pharmacy	0.46	0.47	0.54	0.46	0.46	0.55
	(0.23)	(0.18)		(0.18)	(0.26)	
Perceived days to delivery at private pharmacy	8.52	8.53	1.00	9.71	7.67	0.80
	(12.00)	(12.73)		(14.74)	(9.49)	
Observations	514	312		216	298	

Notes: Columns (1) and (2) display the mean and standard deviation of different covariates at baseline for sample non-attriters and attriters, respectively. Column (3) displays the p-value from a test of equality of means across both groups. Columns (4) and (5) display the mean and standard deviation of different covariates at baseline for treatment and control group within the group of non-attriters surveyed at follow-up. Column (6) displays the p-value from a test of equality of means across both groups within the group of non-attriters surveyed at follow-up.

Table A.4.7: Was a treatment delivered?

	(1)	(2)	(3)	(4)
	Delivered	Explained	Content	Useful
Treatment	0.107***	0.238***	0.304***	0.624
	(0.033)	(0.043)	(0.059)	(0.438)
Constant	0.769***	0.440***	0.379***	7.208***
	(0.025)	(0.033)	(0.049)	(0.379)
Observations	514	514	297	191
R-squared	0.020	0.060	0.083	0.011

Notes: This table displays results from different regressions of measures of treatment delivery on indicators for each of the treatment groups. Column (1) uses an indicator for treatment delivery as an outcome; column (2) uses an indicator for a treatment's being explained; column (3) uses an indicator for whether the participant recalls that the treatment was related to public pharmacies, conditional on receiving it; and column (4) uses a response on a scale from 1 to 10 regarding the usefulness of information, conditional on recalling the content. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.4.8: Balance in covariates between treatment and control group

	(1)	(2)	(3)
Variable	Control	Treatment	p-value $H_0 : (1) = (2)$
Monthly drug expenditure	76.31	76.69	0.94
	(73.54)	(69.97)	
Chronic condition in household	0.57	0.56	0.84
	(0.50)	(0.50)	
Age	45.25	46.32	0.39
	(16.81)	(17.50)	
Education higher than HS	0.54	0.51	0.44
	(0.50)	(0.50)	
Female	0.60	0.63	0.47
	(0.49)	(0.48)	
Public insurance	0.62	0.65	0.37
	(0.49)	(0.48)	
Days with internet per week (1-7)	5.47	5.23	0.23
	(2.71)	(2.84)	
Days with social media per week (1-7)	5.37	5.19	0.37
	(2.79)	(2.91)	
Employed	0.62	0.64	0.53
	(0.49)	(0.48)	
Supports incumbent	0.50	0.51	0.86
	(0.50)	(0.50)	
Voted in previous election	0.73	0.74	0.68
	(0.44)	(0.44)	
Knows public pharmacy	0.61	0.67	0.09
	(0.49)	(0.47)	
Perceived relative price of public pharmacy	0.46	0.46	0.96
	(0.18)	(0.23)	
Perceived days to delivery at private pharmacy	8.80	8.35	0.61
	(12.87)	(11.87)	
Observations	319	507	

Notes: Columns (1) and (2) display the mean and standard deviation of different covariates at baseline for each experimental group. Column (3) displays the p-value from a test of equality of means across the groups.

Table A.4.9: Experimental results for economic outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Panel A - Knowledge about public pharmacies								
	1(Knows about pharmacy)			log(Perceived price)			log(Perceived waiting time)		
Treatment	0.099***	0.069***		-0.117**	-0.094**		0.173	0.188*	
	(0.034)	(0.026)		(0.046)	(0.045)		(0.107)	(0.103)	
Treatment × chronic			0.032			-0.114*			0.134
			(0.033)			(0.061)			(0.140)
Treatment × non-chronic			0.126***			-0.063			0.264*
			(0.042)			(0.065)			(0.151)
Dependent variable at baseline		0.489***	0.488***		0.382***	0.382***		0.397***	0.399***
		(0.039)	(0.039)		(0.049)	(0.049)		(0.068)	(0.068)
Lee bounds	[-0.018, 0.134***]			[-0.236***, -0.020]			[0.049 , 0.189]		
p-value for H_0: $\beta_C = \beta_{NC}$	-	-	0.080	-	-	0.570	-	-	0.531
Mean for control group	0.773	0.773	0.773	9.070	9.070	9.070	1.387	1.387	1.387
Observations	514	514	514	498	491	491	445	425	425
R-squared	0.017	0.474	0.477	0.012	0.197	0.197	0.006	0.181	0.182
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
County FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
	Panel B - Usage of public pharmacies								
	1(Enrolled)			1(Purchased)			Probability of usage		
Treatment	0.018	0.020		0.019	0.023		0.060*	0.054	
	(0.024)	(0.024)		(0.017)	(0.018)		(0.035)	(0.036)	
Treatment × chronic (β_C)			0.032			0.043*			0.085*
			(0.033)			(0.024)			(0.046)
Treatment × non-chronic (β_{NC})			0.002			-0.008			-0.008
			(0.034)			(0.026)			(0.057)
Knows pharmacy at baseline		0.050**	0.050**		0.015	0.015		-0.042	-0.045
		(0.021)	(0.021)		(0.017)	(0.017)		(0.043)	(0.043)
Lee bounds	[0.007, 0.087***]			[0.015, 0.047***]			[0.060 , 0.083]		
p-value for H_0: $\beta_C = \beta_{NC}$	-	-	0.524	-	-	0.155	-	-	0.213
Mean for control group	0.069	0.069	0.069	0.028	0.028	0.028	0.540	0.540	0.540
Observations	514	514	514	514	514	514	387	387	387
R-squared	0.001	0.021	0.100	0.002	0.008	0.067	0.008	0.008	0.057
Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
County FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes

Notes: This table displays cross-sectional estimates using data from the field experiment. In particular, we present results using self-reported indicators about awareness and usage as dependent variables, on the treatment indicator and interactions with and indicator for chronic conditions. Columns 1, 4, and 7 include only a treatment indicator on the right-hand side; columns 2, 5, and 8 include the baseline level of the dependent variable, additional control variables, and county fixed effects; and columns 3, 6, and 9 add an interaction of the treatment indicator with an indicator for whether a member of the consumer household has a chronic condition. The set of control variables includes age and indicators for chronic condition, having completed high school education, female, and public insurance. Outcomes in Panel B either do not have baseline counterparts (which is the case by design of indicators for enrollment and purchase) or were not collected at baseline (which is the case for the probability of usage), so we instead control for knowledge of the public pharmacy at baseline. Reported Lee bounds are computed using only the treatment indicator as a covariate. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix B

Managers and Public Hospital Performance

B.1 Description of the Referral and Counter-Referral System

Other than patients admitted via ER, public hospitals only accept patients referred by other public care centers. Individuals are assigned to a primary care center depending on where they live or work. Referrals to a hospital depend on three main factors: the location of the primary care center and the diagnosis and demographics of the patient. Each Health Service develops detailed referral and counter referral guidelines for all healthcare centers under their territorial scope. Each primary care center can only refer patients following the guidelines defined by the Health Service that supervises them.

Figure B.6.11 illustrates an example of patient referral based on their primary care center. The figure depicts two primary care centers, CESFAM Dra. Haydee López Cassou (in blue with a white diamond marker) and CESFAM Pablo de Rokha (in blue with a white a star marker), which are located in adjacent Health Services. Although individuals in each primary care center might live close to each other, if they require tertiary care they are referred to different hospitals. For most diagnoses, CESFAM Dra. Haydee López Cassou refers their adult patients to Hospital Barros Luco (in red with a white a cross marker) and CESFAM Pablo de Rokha refers them to "Hospital Sótero del Río" (in red with a white H marker).

Table B.6.1 shows an example of referral guidelines from different primary care centers to public hospitals in two Health Services. Primary care centers in columns (1)-(2) and (3)-(4) are in two different Health Services: Metropolitano Norte and Metropolitano Oriente, respectively. The numbers in the table are the hospital to which patients are referred. The example shows that referrals depend exclusively on the primary care center and the diagnosis and demographics of the patient. For example, a medical oncology patient older than 15 in CESFAM Colina is referred to "Instituto Nacional del Cáncer Dr. Caupolicán Pardo Correa."

To empirically assess compliance with the referral guidelines, we focus on a sample of patients with public insurance who were discharged (dead or alive) at any point during the year 2004 and who were not admitted into the hospital via ER. In this sample, we classify patients into cells defined by patients' county of residence, age group (less than 1, between 1 and 15, and more than 15) and diagnosis (as reported by the hospital from which they are discharged). If the guidelines are strictly followed, we should expect all patients within a cell to attend the same hospital. To visually evaluate this, Figure B.6.12 plots a spikeline with the share of patients in each cell who are discharged exclusively from one hospital; more than 80% of patients within a cell are discharged from the same hospital. Importantly, the fact that patients within a cell are being discharged from different hospitals does not necessarily constitute evidence of non-compliance with the referral and counter referral guidelines. In our case, this may reflect censorship due to the fact that we do not observe the diagnosis at the primary care center, but only at the hospital. Likewise, this could be explained by the fact that we only observe patients' home address, but they could have used their work address to register in the health system. Finally, there might also be measurement error in the address and age of patients.

B.2 Managers matter for hospital performance

In this appendix, we explore the rotation of CEOs across hospitals to study the extent to which CEOs affect hospital quality. Specifically, we follow the approach used by Fenizia (2022) and exploit the rotation of CEOs across hospitals to estimate the following model:

$$\text{Ln(death rate)}_{ht} = \alpha_h + \psi_{M(h,t)} + \gamma_t + X'_{ht}\Delta + u_{ht}, \tag{B.1}$$

where α_h are hospital fixed effects that capture time-invariant characteristics of the hospital (e.g., size and the type of procedures performed there), and $\psi_{M(h,t)}$ are CEO fixed effects, which capture managerial talent (specific to a given CEO) and are assumed to be portable across hospitals. We also include time fixed effects γ_t to capture seasonal shocks to patients' health and health provision as well as case mix controls, X_{ht}, to account for differences in patients' demographics (age and sex) and socioeconomic status (proxied by type of insurance).

For estimation, we first identify the set of hospitals that are connected by CEOs' mobility (Abowd, Kramarz, and Margolis, 1999; Card et al., 2013) and define our main estimation sample, which consists of 789 CEOs, 113 hospitals, and 19 connected sets created by 86 movers. Then, we estimate the model via constrained OLS and recover CEO fixed effects that can be compared *within* connected sets.

It is worth noticing that models with additive hospital and CEO components may raise some concerns. One may worry, for instance, that CEOs are assigned to hospitals on the basis of unobserved factors that determine their comparative advantage. It could also be that manager rotation is correlated with hospital-specific trends. In the next subsection we empirically assess these concerns as in Card et al. (2013) and Fenizia (2022). All in all, the evidence suggests that the two-way fixed effects model fits the data well, and match effects, if any, are small.

Threats to Identification

We follow Card et al. (2013) and Fenizia (2022) to assess the two main threats to the identification of $\hat{\psi}_{M(h,t)}$. The first concern is that CEO mobility is endogenous due to a systematic relation with hospital-specific trends—for example, if good CEOs are rotating to hospitals that are improving their quality over time. This pattern would overestimate our CEO fixed effects. Relatedly, one might worry that CEOs move to a new hospital due to transitory productivity shocks in that hospital. This would be the case, for instance, if a given hospital performs poorly in a given period and, in response, makes an extra effort to hire a good manager. To assess this concern, we exploit the rotations of CEOs in an event study framework. Specifically, we calculate the difference between the incumbent and the incoming CEO (hereafter, Δ CEO FE) and classify CEO transitions into terciles. Intuitively, the classification allows us to distinguish whether the new CEO implies an average increase, a small change, or an average decrease in manager quality.

Panel (a) in Figure B.6.13 plots the effect of CEO transitions on residualized death rates for each Δ CEO FE tercile. Several points are worth noting about this figure. First, hospitals with an event in the first tercile observe a significant decline in death rates after the CEO changes, and the opposite is true for events in the third tercile. In both cases, the effects persist over time. Moreover, we find no effect on hospital quality for Δ CEO FE in the second tercile, in which changes in CEO quality are small. A second observation is that hospitals that hire a good or bad incoming CEO (relative to the incumbent) are not on different trends, and that turnovers do not seem to correlate with pre-trends of hospital performance. Before a CEO turnover, hospitals that face a CEO move exhibit a trend similar to those that do not, consistent with evidence presented in Figure 2.2. In sum, we think that these event studies should ameliorate concerns regarding endogenous mobility.

The second threat to the identification of manager fixed effects comes from the potential existence of match effects between CEOs and hospitals; this dimension is neglected in the log model by the additive separability between CEO and hospital effects. Different CEOs may have different effects on hospital quality, depending on the value of their match component. If CEOs sort into hospitals in which they have a comparative advantage, this effect would be captured by the error term and would bias our estimates. To examine whether this concern is valid we consider two pieces of evidence. First, in column (6) of Table 2.1, we report a saturated version of Equation B.1, in which we include CEO-by-hospital fixed effects. If the match component is sizable, this model should have a better fit than that in column (4). We find that the adjusted R^2 increases from 0.69 to 0.72 after including manager-by-hospital fixed effects—a rather modest change in model fit. We further examine to what extent the model is overlooking match effects by analyzing whether the mean residuals are abnormally high or low for a given pair of hospital and CEO. With this in mind, we divide the estimated manager and hospital effects into quartiles and compute the mean residual for each pair. Results are depicted in Panel (b) in Figure B.6.13. We find that all residuals are small and lower than 0.05 in absolute value. A final piece of evidence comes from the symmetry of the effects depicted in Panel (a) in Figure B.6.13. Hospitals that move from a good CEO (in the first tercile) to a bad CEO (third tercile) face an opposite and symmetric effect to that of moving from a bad CEO to a good CEO, which would be implied by the lack of match effects in the model. All in all, the evidence suggests that the two-way fixed effects model fits the data well, and match effects, if any, are small.

Variance Decomposition

In this subsection we perform the variance decomposition. Following Equation B.1, the variance of log death rates can be decomposed as:

$$\begin{aligned}\mathbb{V}(\text{Ln(death rate)}_{ht}) =&\mathbb{V}(\alpha_h) + \mathbb{V}(\psi_{M(h,t)}) + \mathbb{V}(x'_{ht}\beta) + 2\mathbb{C}(\alpha_h, \psi_{M(h,t)}) \\ &+2\mathbb{C}(\alpha_h, x'_{ht}\beta) + 2\mathbb{C}(\psi_{M(h,t)}, x'_{ht}\beta) + \mathbb{V}(u_{ht}),\end{aligned} \tag{B.2}$$

where x_{ht} includes patients' demographics (age and sex), socioeconomic status (proxied by type of insurance) and time effects. Table B.6.3 presents the magnitude of each term in Equation B.2, estimated within the largest connected set.[1] Since sampling error could bias the estimates in the presence of limited mobility, we correct the estimates following the procedure of Andrews et al. (2008). We find that most the variance in death rates is explained by patients' characteristics and time effects (76%). Manager fixed effects explain around 26% of the variance in death rates, which is about 70% of the permanent component associated with different hospitals which explains 36% of the variance in death rates. Our results also show that the (bias-corrected) covariance between CEO and hospital effects is negative (-10%), which implies that the most talented CEOs work at least productive hospitals (i.e., there is negative assortative matching).

B.3 CMS Risk Adjustment

To ease selection concerns, we follow the UK's National Health Service (NHS) (e.g., Health and Centre, 2015) and construct a "risk-adjusted mortality rate" that divides the actual hospital-level death rate by the death rate predicted based on the observable characteristics of hospitals' patients. This variable should be interpreted such that an increase (decrease) from one means a larger (smaller) death rate than predicted based on hospital case mix.

The prediction is built following the procedure described in Ash et al. (2012), which the Centers for Medicare and Medicaid Services (CMS) use in the United States. First, we focus on a sample of 5,740,496 patients between 2001 and 2004 (before reform adoption). These patients constitute the universe of discharges in the country. For them, we fit a logit model

[1] The largest connected set in our setting contains 3,276 observations: 322 CEOs, 41 hospitals, and 46 movers.

in which death is the dependent variable and different sets of patients' characteristics are the independent variables. Our preferred model includes the following set of covariates: gender; eight age buckets (< 30, 30 – 49, 50 – 59, 60 – 69, 70 – 79, 80 – 89, and > 89); type of health insurance (private or one of 5 categories within public insurance that depend on income); and the 31 categories of the enhanced Elixhauser comorbidity index (Elixhauser et al., 1998; Quan et al., 2005). Then, we predict the probability of death for each patient, which is a variable we use to construct the predicted death rate at hospital level.

B.4 No Differential Effects of Performance Pay

In this appendix, we empirically examine whether CEOs' scores on their performance pay measure predicts better managerial performance in the hospital. We define a dummy variable that takes value 1 if the manager was above the performance score median and 0 otherwise. We interact this variable with introduction of the reform and study the impact of the reform for managers with high and low scores separately.

Appendix Table B.6.5 displays the results. Since we do not observe performance scores for all managers who took over after the reform, we miss several observations. For this reason, in column (1) we report the impact of the reform in the sample for which we have data. Importantly, we find the same effect as when we use the whole sample. In column (2), we report the results of the reform for managers with high and low scores. We find that both estimates are almost identical. As posited above, this is evidence that performance pay did not have any effect on manager performance.

B.5 Effect of the Selection Reform on CEO Wages

In this appendix, we study the reform's effect on hospital CEOs' wages. We leverage the gradual adoption of the reform across public hospitals and estimate an event study specification on the position's wage. An important caveat is that the wage data panel starts in January 2014, after which only three hospitals adopted the selection reform for the first time. Fortunately, we also have data for December 2011-2013, which gives us a larger number of events. For this reason, we also estimate an event study using data only for December, between 2011 and 2019.

Panel (a) in Figure B.6.15 presents the results using quarterly data starting on 2014. Although the estimates are noisy due to the small number of events, the estimate is stable and the average quarterly wage increase in the 5 quarters post-adoption is 33%. We also do not find evidence of pre-trends, which means that hospitals that adopt the reform during a given period are not on a wage trend that differs from those that do not. Panel (b) presents estimates using monthly data for each December, starting in 2011. In both cases, standard errors are clustered at hospital level. We find quantitative and qualitatively similar results.

It is important to note that the results of this exercise reflect the change in the position's pay, and therefore is a composite of two effects. On the one hand, there are mechanical changes in pay due to changes in the manager's characteristics. For example, in the public sector, there are tenure bonuses that increase with experience. On the other hand, there is an increase in the position's base wage. Since our wage data follow the position and not the individuals over time, we cannot separate the effects.

B.6 Additional Figures and Tables

Figure B.6.1: Share of medical beds provided by public hospitals in OECD economies

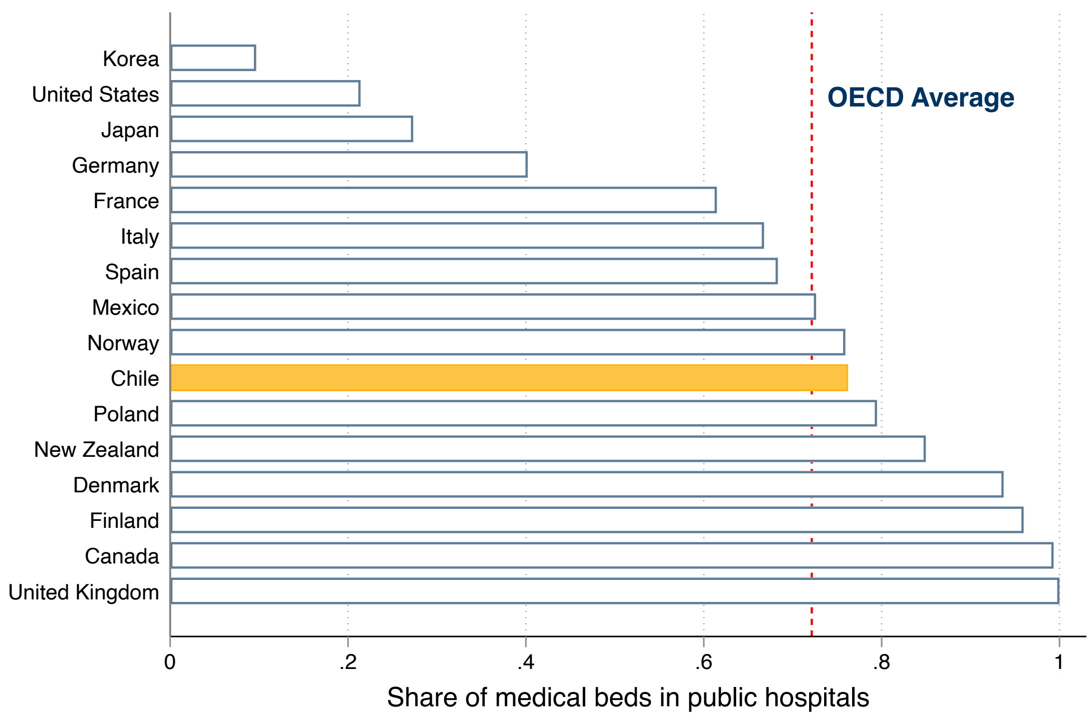

Notes: This figure displays the share of medical beds provided by public hospitals in a set of selected OECD countries in 2019. The dashed red line represents the average share in all OECD countries. The share is computed as the ratio between the total number of hospital beds in publicly owned hospitals and the total hospital beds in the country. Both variables are reported in OECD (2022a).

Figure B.6.2: Health Services are distributed geographically

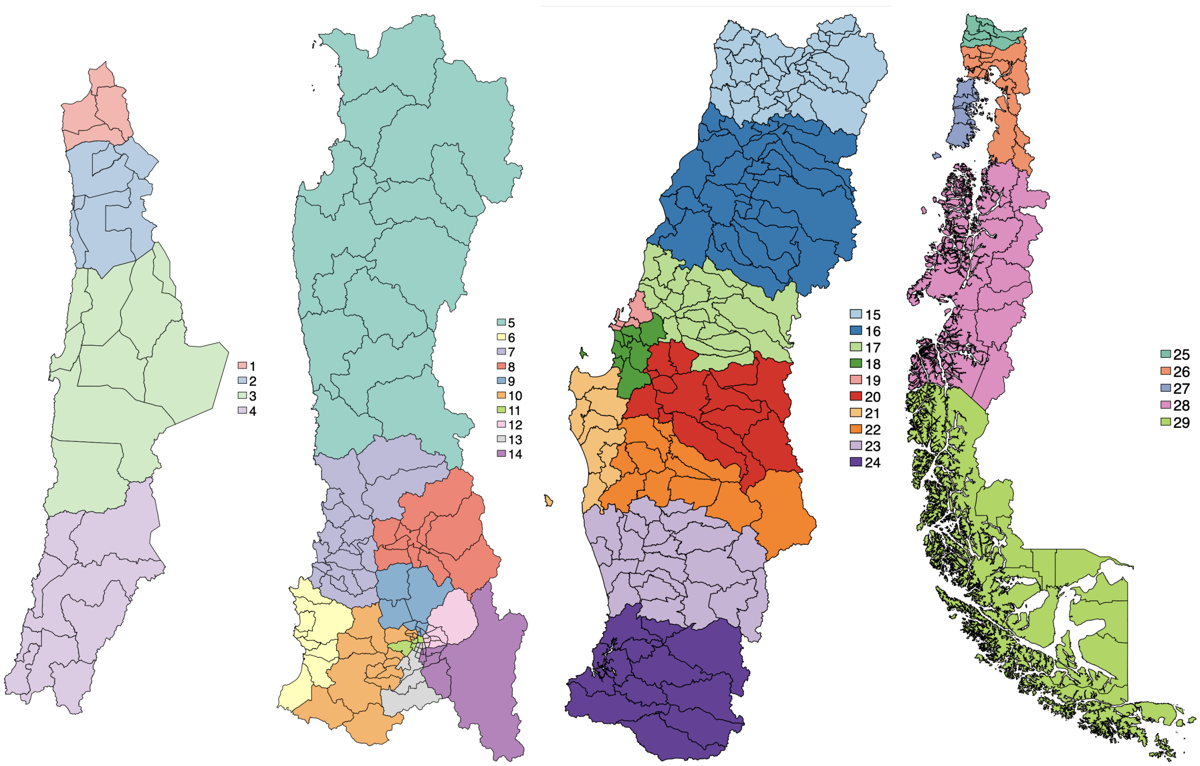

(a) North (b) North-Center (c) Center-South (d) South

Notes: This figure shows the geographic distribution of the 29 Health Services in Chile. Each Health Service is responsible to oversee public health providers in the municipalities in their territory. Colors represent different Health Services and black lines represent municipal borders.

Figure B.6.3: Selection process after the recruitment reform

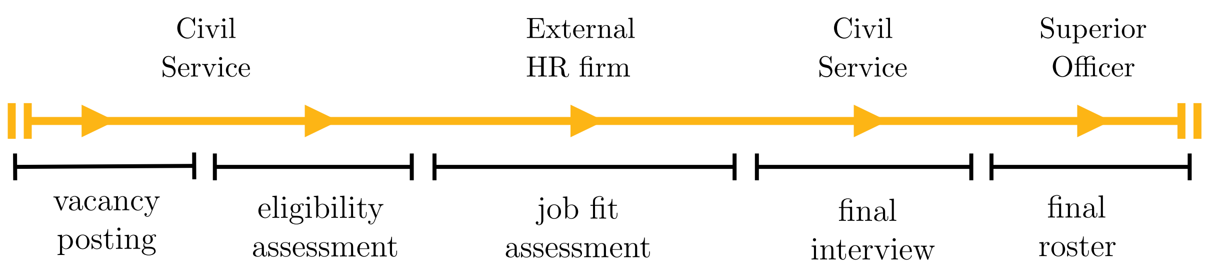

Notes: This figure illustrates the selection process for senior executives positions when the selection reform has been adopted. The job call starts with the position posted online on the Civil Service's website and in a newspaper with national circulation. In some cases, the Civil Service may also hire headhunters to widen the pool of applicants. After the job posting closes, an external HR firm evaluates each individual's job trajectory according to the job profile. They also assess motivation and overall competencies. The consultant gives every applicant a grade based on an objective rubric and provides a short list to the Civil Service. In the next phase, a committee consisting of representatives of the Civil Service and the Ministry in which the position is based interviews the remaining candidates and selects a short list of three individuals based on objective criteria. Finally, the superior officer appoints the winning candidate from the final roster with complete discretion.

Figure B.6.4: Gradual adoption of the reform by public hospitals

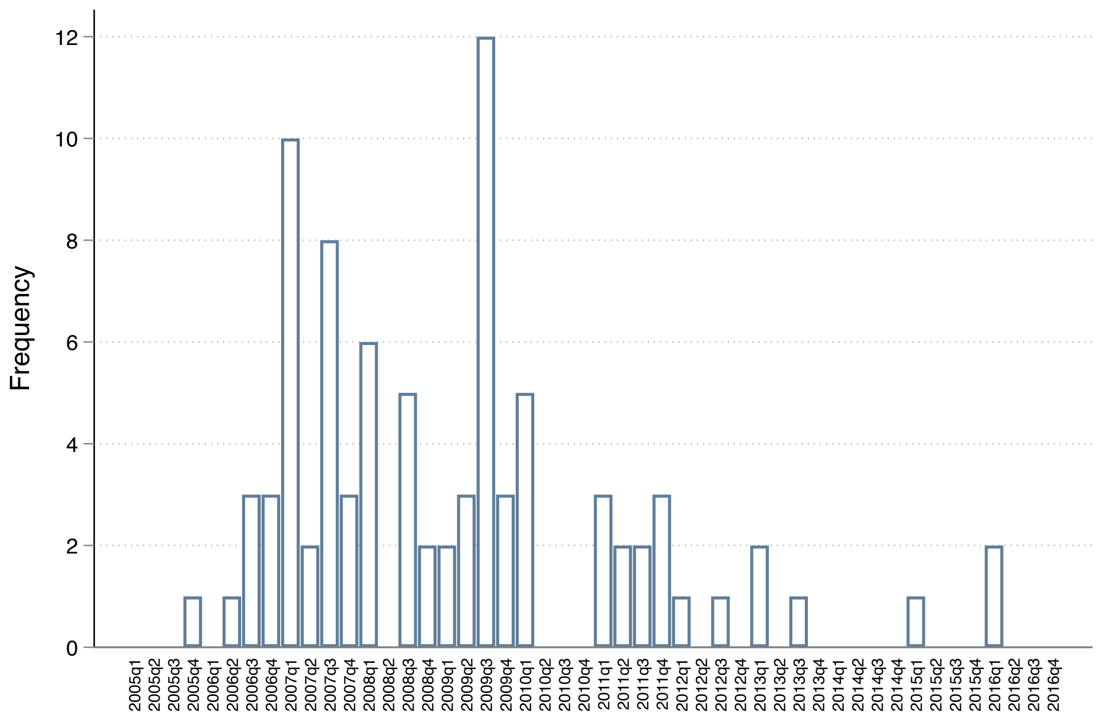

Notes: This figure displays adoption of the selection reform for CEOs in public hospitals. An observation represents a hospital that adopts the selection reform for the first time. After a hospital adopts the process, all future CEOs in that hospital have to be appointed using the new selection system.

Figure B.6.5: Dynamic effects of the reform using alternative models

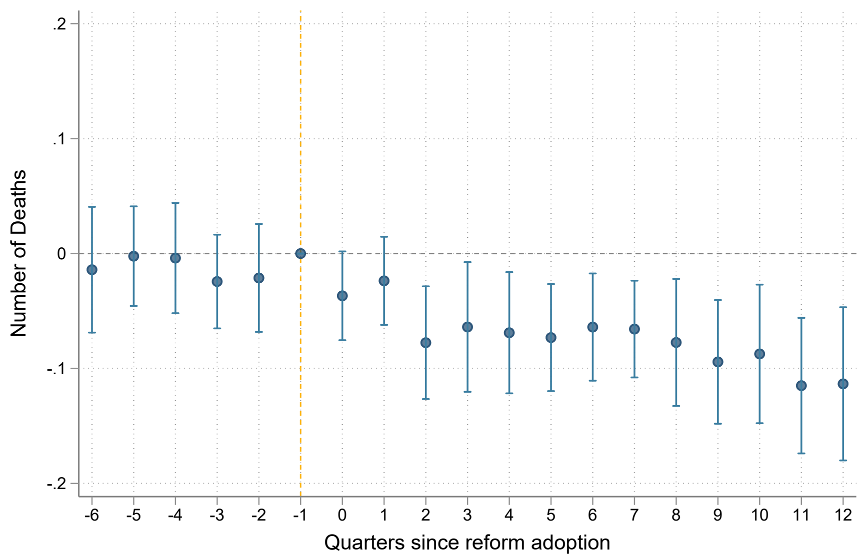

(a) Poisson Model

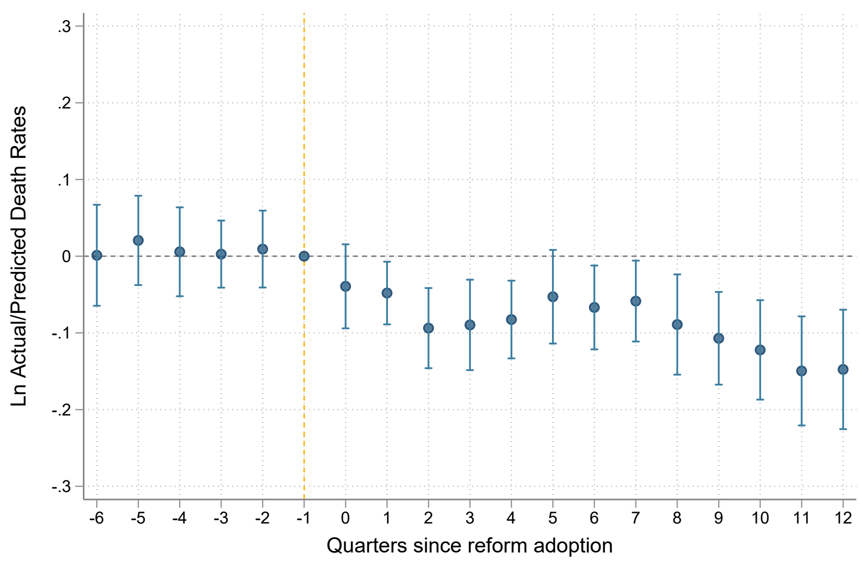

(b) Risk-Adjusted Mortality

Notes: This figure presents the reform's effect on alternative outcome variables. Panel A reports estimates obtained from a dynamic Poisson regression of death counts. Panel B reports estimates and confidence intervals obtained from a two-way fixed-effects OLS regression of logged risk-adjusted death rates. The average death rate is predicted by patient-level characteristics using a logit model for deaths (for details, see Appendix B.3). We define the risk-adjusted death rate as the actual hospital-level death rate divided by the predicted rate. Each dot corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure B.6.6: Alternative event study models and estimation methods

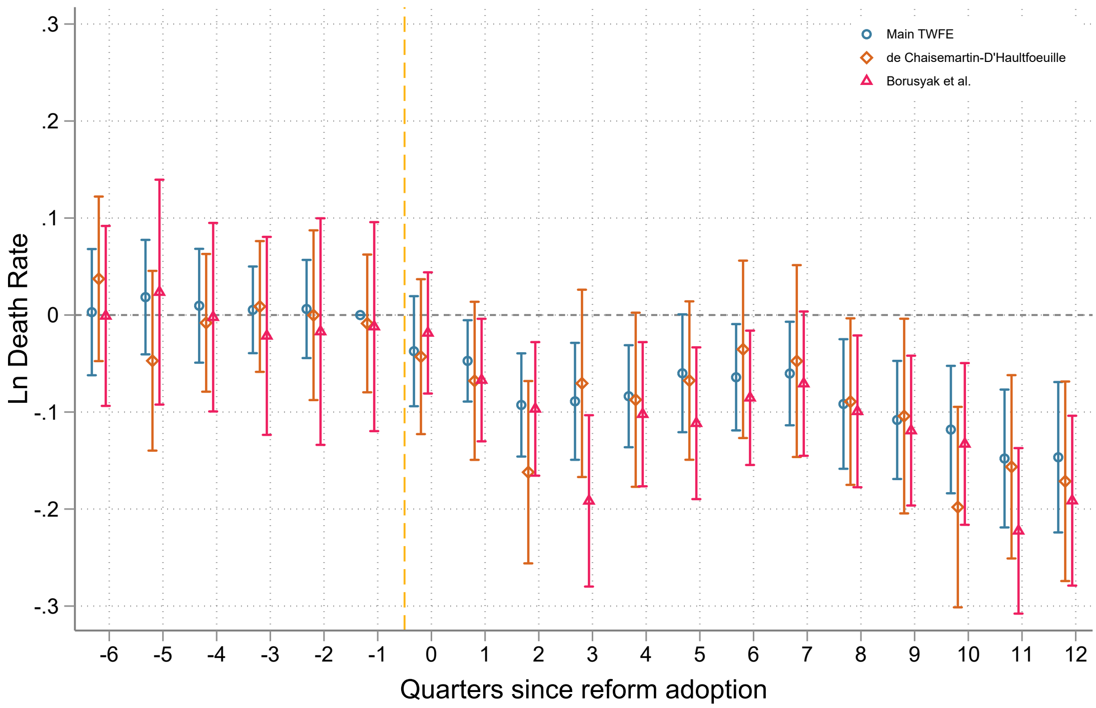

Notes: This figure plots the estimates and confidence intervals obtained from different event study models and estimation methods. The main event-study results using a two-way fixed-effects (TWFE) regression of logged death rates (see Equation 2.2) are presented under the label "Main TWFE" (in blue with circle markers). For comparison, we overlay the results obtained using the models suggested by De Chaisemartin and d'Haultfoeuille (2020) (in orange with diamond markers) and Borusyak et al. (2022a) (in red with triangle markers), which are robust to treatment effect heterogeneity. Each dot, diamond, and triangle marker corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals.

Figure B.6.7: Effects on predicted mortality

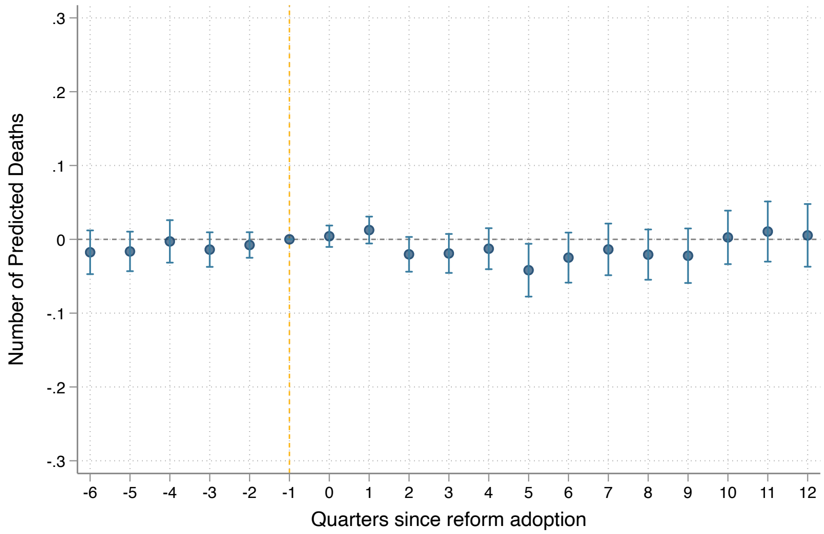

(a) Predicted Number of Deaths

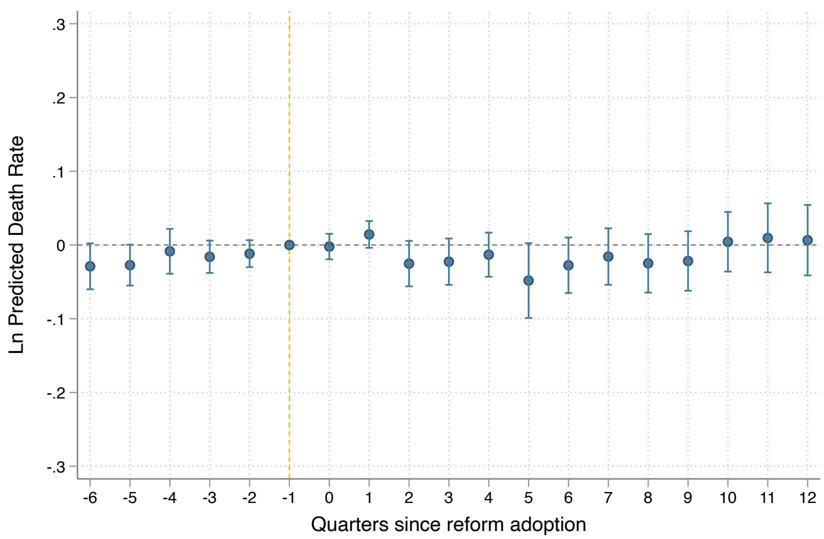

(b) Predicted Death Rate

Notes: This figure presents event study evidence on the null effects of the reform on mortality when predicted based on patients' risk-scores. For this exercise, we fit a logit model of deaths on patients' demographics and diagnoses (for details, see Appendix B.3). Then, we predict the probability of death for each patient, and use these predictions (i.e., patient level risk scores) to construct the hospital-level predicted death rate and number of deaths. Panel A reports estimates obtained from a dynamic Poisson regression of predicted death counts. Panel B reports estimates obtained from a dynamic two-way fixed-effects OLS regression of logged predicted death rates. Each dot corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure B.6.8: Effect of the reform on hospital performance (year-level aggregates)

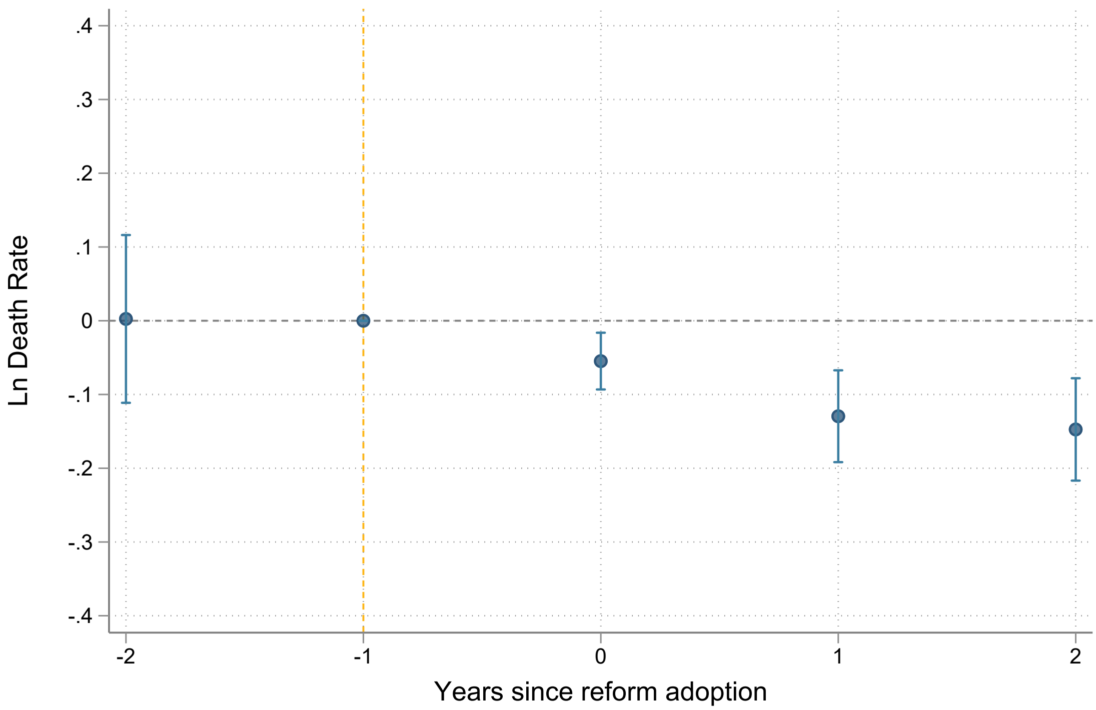

Notes: This figure presents event study evidence of the reform's effect on death rates when the outcome is logged hospital death rates at year level. Each dot corresponds to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. Dashed yellow lines represent the omitted coefficient. Standard errors are clustered at hospital level.

Figure B.6.9: Creation of postgraduate programs in health management

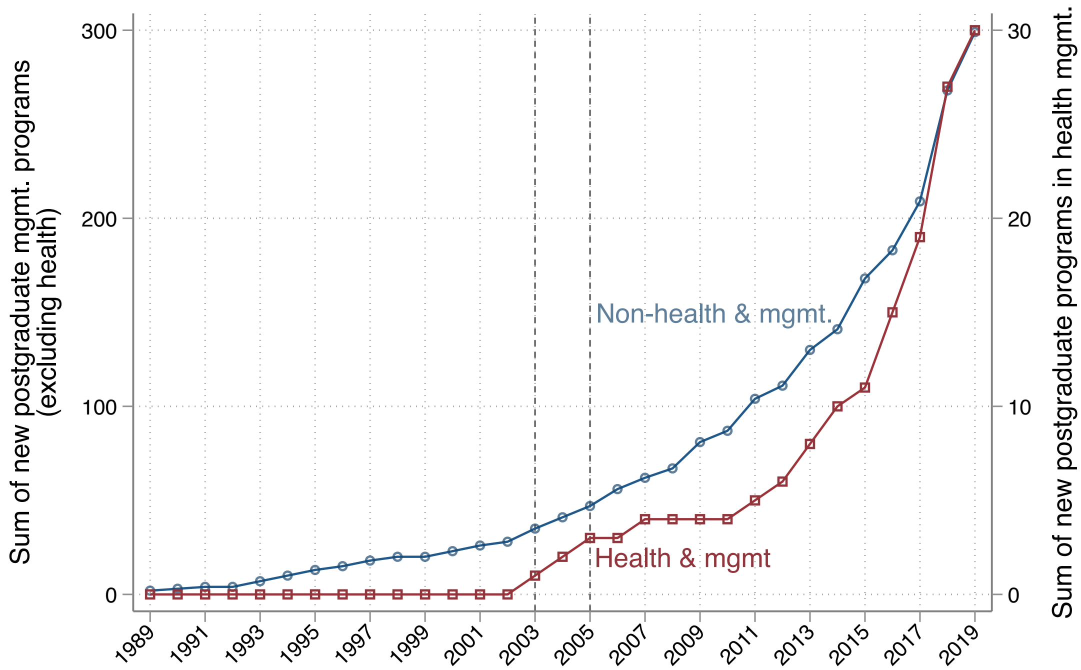

Notes: This figure shows the cumulative number of postgraduate management programs (diplomas and master's) by date of creation. The blue circles depict all management postgraduate degrees, excluding those related to health; corresponding frequencies are displayed in the left y-axis. The red squares depict new postgraduate degrees that include both management *and* health in their descriptions; corresponding frequencies displayed in the right y-axis. Dashed gray lines indicate the years when Law Nº 19,882 (which created the new selection system in the country) was enacted and when the first hospital adopted the new selection system. We use data from programs actively running in 2019, reported by the Consejo Nacional de Educación (https://www.cned.cl/bases-de-datos).

Figure B.6.10: Share of total CEO wage explained by the reform's bonus

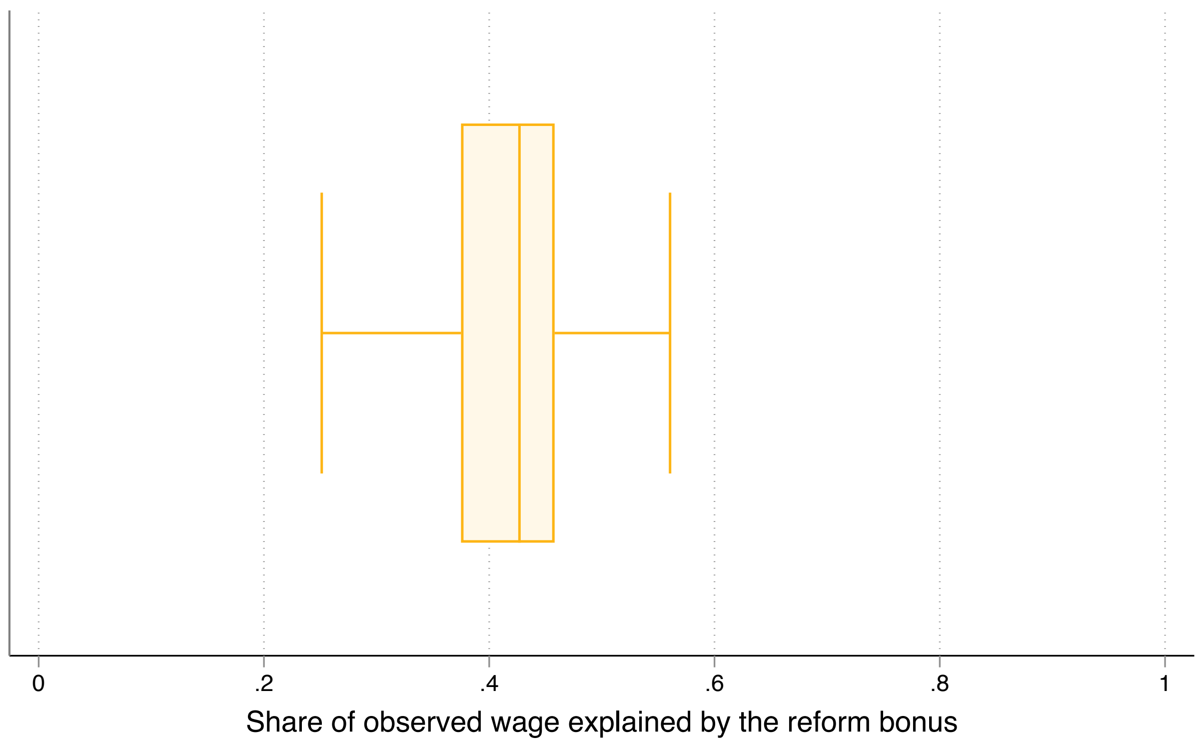

Notes: This figure displays a box plot of the share of the CEO wage explained by the reform's bonus. The sample consists of all CEO positions appointed using the selection reform between 2014 and 2019, which are the dates for which monthly data are available. The average wage share explained by the reform's bonus is 43%. The 25th and 75th percentiles are 37% and 46%, respectively. Figure excludes outside values.

Figure B.6.11: Examples of referral from primary care centers

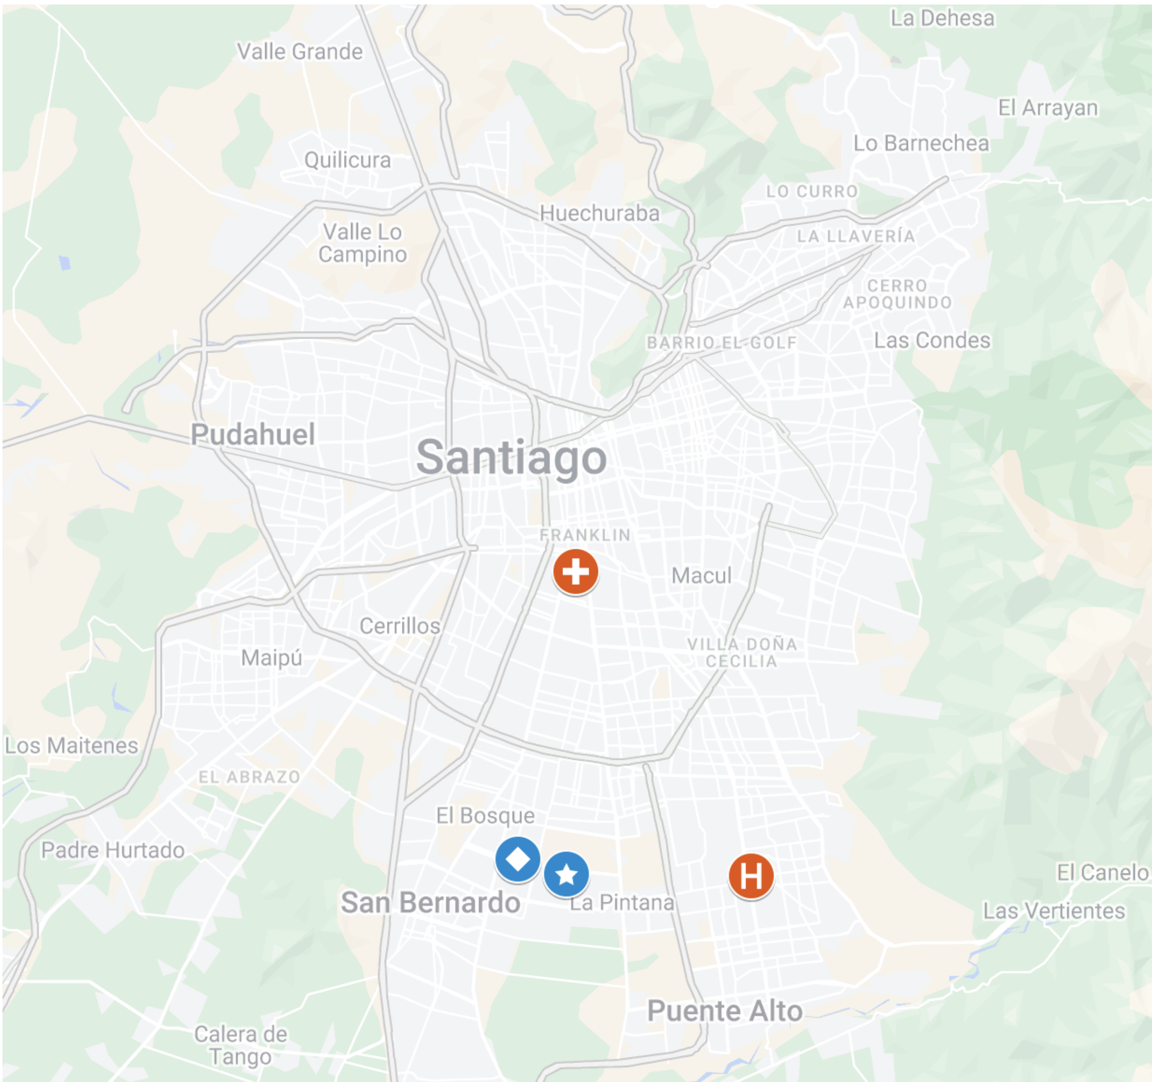

Notes: This figure illustrates an example of patient referral based on their primary care center. The figure depicts two primary care centers, CESFAM Dra. Haydee López Cassou (in blue with a white diamond marker) and CESFAM Pablo de Rokha (in blue with a white star marker), which are located in adjacent Health Services. Although individuals in each primary care center might live close to each other, if they require tertiary care they will be referred to different hospitals. For most diagnoses, CESFAM Dra. Haydee López Cassou refers their patients to Hospital Barros Luco (in red with a white cross marker). Patients from CESFAM Pablo de Rokha are referred to Hospital Sótero del Río (in red with a white H marker). Referrals depend exclusively on the location where the individual is enrolled, her diagnosis, and her demographics. Table B.6.1 shows an example of referrals to different public hospitals within the same Health Service based on the patient's diagnosis and demographics.

Figure B.6.12: Empirical test of patient selection

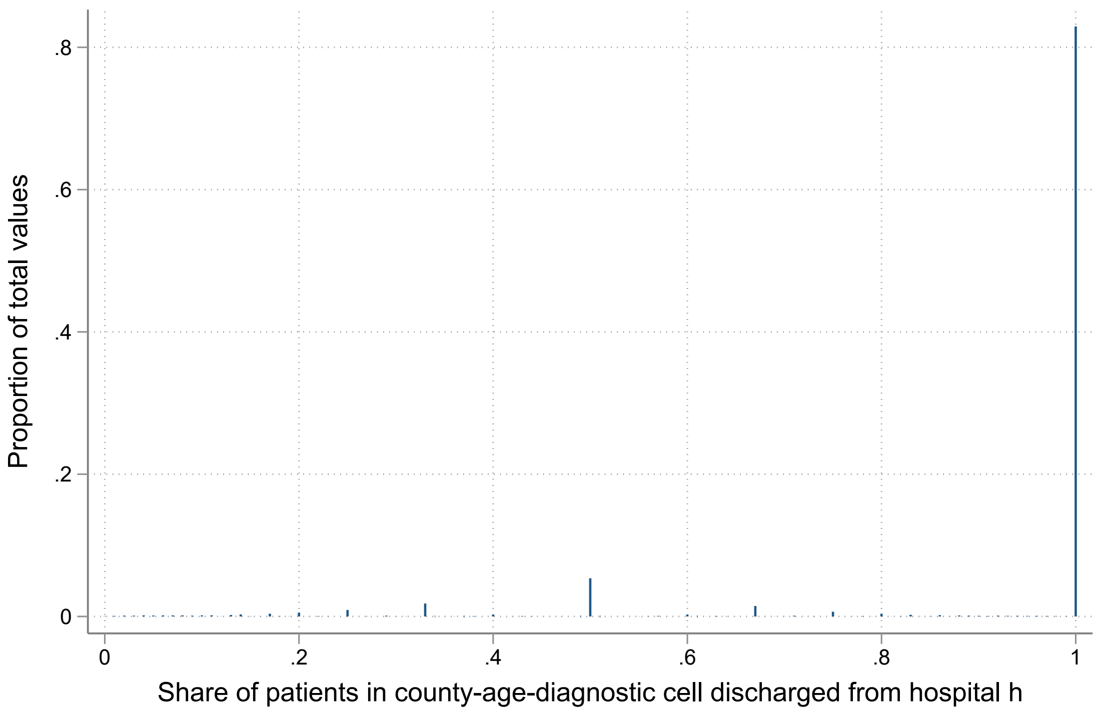

Notes: This figure plots a spikeline with the share of patients in each cell who are discharged exclusively from one hospital. A cell is defined by the patient's county of residence, age group (less than 1 year, between 1 and 15 year, and more than 15 year) and diagnosis (as reported by the hospital from which they are discharged). If referral guidelines are strictly followed, we should expect all patients within a cell to attend the same hospital. However, in our data patients within the same cell could be discharged from different hospitals due to the fact that we do not observe the diagnosis at the primary care center, but only at the hospital. Likewise, it may be due to the fact that we only observe patients' home address, but they could have used their work address to register in the health system.

Figure B.6.13: Threats to the identification of managerial talent

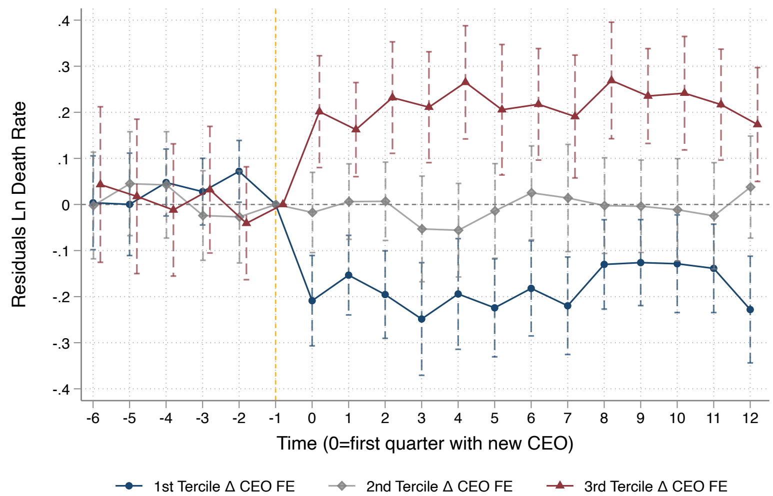

(a) Switchers

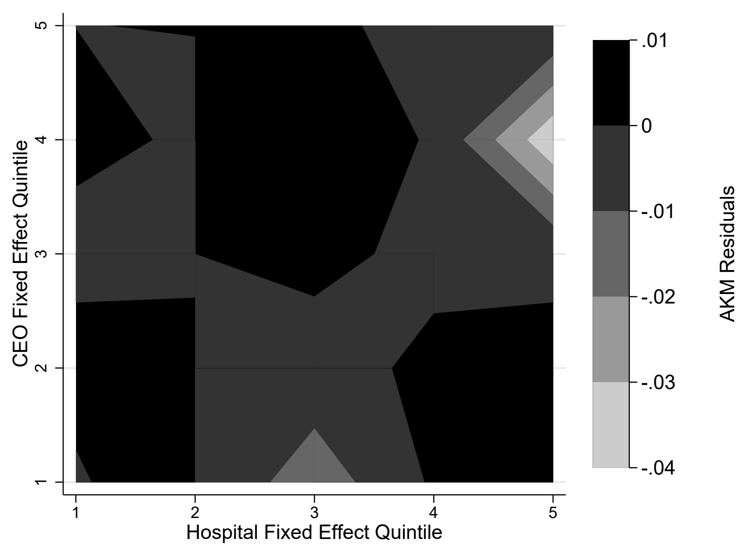

(b) Residuals

Notes: This figure presents evidence against potential endogenous mobility of managers and in favor of the additive separability assumption between hospital and manager components. Panel A plots the mean (residualized) log death rate against event time (relative to change in CEO events). The figure plots three types of leadership transitions, classified by terciles of the change in managerial ability: (1) an overall increase (in blue with dot markers), (2) an overall decrease (in red with triangle markers), and (3) no significant change (in gray with diamond markers). Each dot, triangle, and diamond marker correspond to an estimated coefficient, and vertical lines indicate the corresponding 95% confidence intervals. Panel B shows mean residuals from model B.1 with cells defined by quintiles of estimated manager effect, interacted with quintiles of estimated hospital effect.

Figure B.6.14: Differential effects of CEO transitions

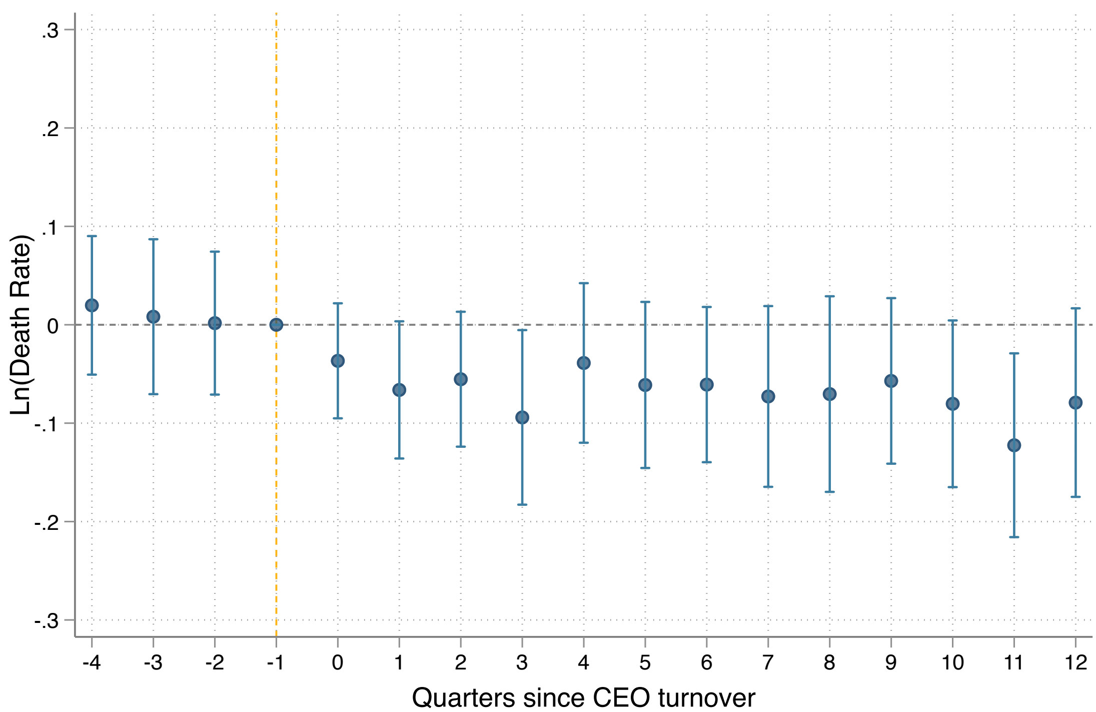

Notes: This figure presents the coefficients of the stacked event study specification in Equation 2.3, in which an event is a transition from a CEO without management studies to a CEO with management studies. For each transition event, we define a time window around it and a control group of hospitals with no transitions in the time window. We define a set of valid events as those that are balanced in the time window and do not overlap with another transition in the pre-period within the time window. We also exclude transitions associated with the first time a CEO was appointed after the selection reform was adopted by a given hospital. In total, there are 94 valid CEO transitions, as described in Appendix Table B.6.4. The dependent variable is the death rate at hospital level in a given quarter. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95% confidence intervals. Standard errors are clustered at hospital level.

Figure B.6.15: Impact of recruitment reform on wages

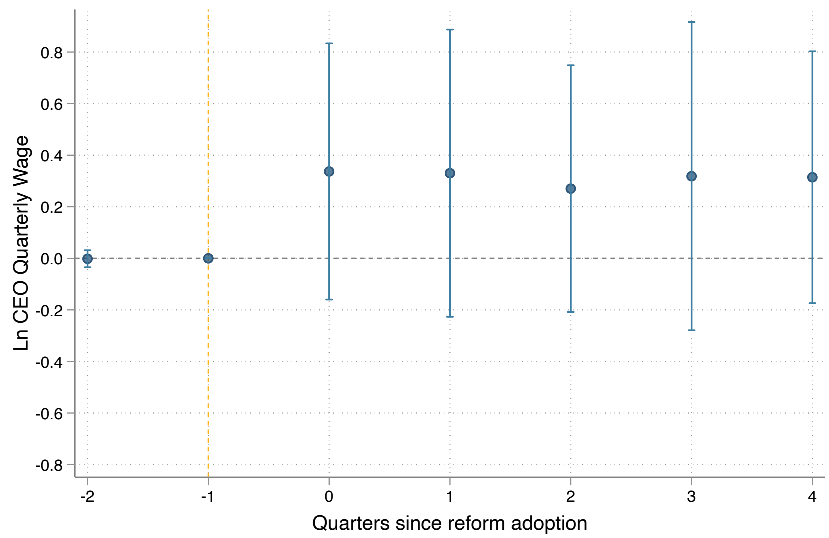

(a) 2014-2019: Quarterly

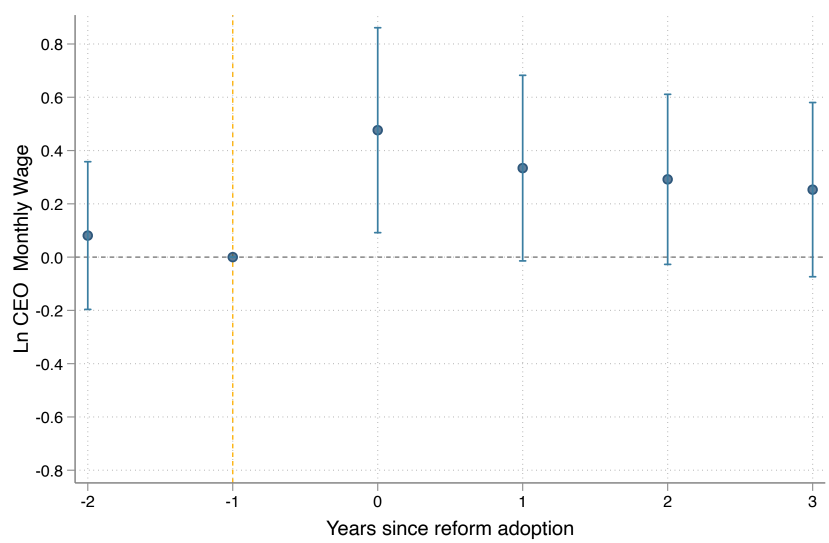

(b) 2011-2019: December only

Notes: This figure presents the impact of the reform on hospital CEO wages. The empirical design leverages the gradual adoption of the selection reform across hospitals. Panel A presents the results using quarterly panel data between 2014 and 2019. Although the estimates are noisy due to the small number of events, the estimate is stable at around 33%. Panel B uses data for December between 2011 and 2019, which allows us to leverage a larger number of events. Regression controls include age and a dummy that indicates whether the individual is a doctor, which affects pay in the public sector. Dots indicate estimated coefficients and vertical lines indicate the corresponding 95% confidence intervals. Standard errors are clustered at hospital level.

Table B.6.1: Referral guidelines example

Health Service Name	*Metropolitano Norte*		*Metropolitano Oriente*	
Primary Care	CESFAM Colina (1)	CESFAM Esmeralda (2)	CESFAM Aguilucho (3)	CESFAM La Faena (4)
Pediatrics				
Pediatric respiratory diseases	2	2	4	4
Internal Medicine				
Cardiology	1	1	5	4
Medical Oncology				
<15 years	2	2	7	7
>15 years	3	3	5	5
General Surgery				
Thoracic Surgery	3	3	6	6

1: Complejo Hospitalario San José
2: Hospital Clínico De Niños Roberto Del Río
3: Instituto Nacional Del Cáncer Dr. Caupolicán Pardo Correa
4: Centro de Referencia de Salud Cordillera Oriente
5: Hospital Del Salvador
6: Instituto Nacional del Torax
7: Hospital de Niños Dr. Luis Calvo Mackenna

Notes: This table illustrates the referral guidelines from primary public care to public hospitals. Referrals depend on the primary care center and the demographics of the patient. Columns (1)-(2) and (3)-(4) are in two different Health Services, Metropolitano Norte and Metropolitano Oriente, respectively. Numbers represent the hospital to which the patient is referred. For example, a patient for medical oncology older than 15 years in CESFAM Colina is referred to the Instituto Nacional del Cáncer Dr. Caupolicán Pardo Correa.

Table B.6.2: Impact on risk-adjusted mortality measures

	Death Rate Ln Actual/Predicted		
	(1)	(2)	(3)
1 if reform adopted in hospital	-0.080***	-0.081***	-0.080***
	(0.022)	(0.022)	(0.022)
Observations	8,104	8,104	8,104
Time FE	Yes	Yes	Yes
Hospital FE	Yes	Yes	Yes
Patient Demographics	Yes	Yes	Yes
Type of Insurance	No	Yes	No
Enhanced Elixhauser Comorbidity Index	No	No	Yes
Pseudo-R^2 Logit	0.147	0.158	0.176
# of Hospitals	181	181	181
Mean Dep. Variable	0.787	0.722	0.747

Notes: This table presents our estimates of the impact of the selection reform on risk-adjusted death rates. The estimates are from the staggered DiD specification in Equation 2.1. We define the risk-adjusted death rate as the actual hospital-level death rate divided by the average death rate as predicted by different patient-level characteristics used to fit a logit model for deaths. See Appendix B.3 for details. Standard errors are displayed in parentheses and are clustered at hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.6.3: Hospital performance variance decomposition

	Component	Share of Total
	(1)	(2)
$\mathbb{V}$(Log Death Rate)	0.526	100%
$\mathbb{V}$(Manager)	0.139	26%
$\mathbb{V}$(Hospital)	0.193	36%
$\mathbb{V}(x'_{ht}\beta)$	0.403	76%
$2\mathbb{C}$(Manager, Hospital)	-0.055	-10%
$2\mathbb{C}(x'_{ht}\beta,$ Manager + Hospital$)$	-0.001	-0.00%
$\mathbb{V}$(Residual)	-0.149	-28%

Notes: This table reports bias-corrected variances and covariances estimated on the largest connected set following Andrews et al. (2008). Hospitals and managers' fixed effects are estimated from Equation B.1 in the set of hospitals connected by managers' mobility (Abowd et al., 1999; Card et al., 2013). $x'_{ht}\beta$ includes patients' demographics (age and sex), socioeconomic status (proxied by type of insurance) and time effects. For details see appendix B.2.

Table B.6.4: CEO transitions according to management studies

	Current CEO has:			*Total*
Previous CEO had:	Non-Mgmt. Studies (1)	Mgmt. Studies (2)	No Data (3)	
Non-Mgmt. Studies	431	94	5	530
Mgmt. Studies	95	66	4	165
No Data	31	4	4	39
Total	557	164	13	734

Notes: This table presents the number of CEO transitions according to the characteristics of the incumbent and incoming manager in terms of management studies (mgmt. studies). We only consider transitions for which there is a time window of 4 periods before and 12 periods after the transition, and do not overlap with another transition in the pre-period within the time window.

Table B.6.5: No differential effects according to performance pay scores

	Ln Death (%) (1)	Ln Death (%) (2)
Reform	-0.087***	
	(0.028)	
Reform & High Score		-0.086**
		(0.033)
Reform & Low Score		-0.089**
		(0.036)
Observations	7,670	7,670
Time FE	Yes	Yes
Hospital FE	Yes	Yes
Case Mix Controls	Yes	Yes
# of Hospitals	181	181
Mean Dep. Variable	2.61	2.61
p-value *High Score* = *Low Score*		0.94

Notes: This table examines differential effects of the recruitment reform depending on the CEO's average performance score. Their performance score is measured according to their performance contract. In column (1), we replicate the estimation of Equation 2.1 in the subsample for which we have performance scores. Column (2) interacts the reform with whether the CEO scored above or below the median in the performance score outcome. Standard errors are displayed in parentheses and are clustered at hospital level. *** p<0.01, ** p<0.05, * p<0.1.

Table B.6.6: Correlation between CEO fixed effect and manager characteristics

	CEO Fixed Effect				
	(1)	(2)	(3)	(4)	(5)
Female	-0.068*	-0.065*	-0.071*	-0.054	-0.052
	(0.037)	(0.036)	(0.036)	(0.035)	(0.035)
Age	0.166***	0.163***	0.163***	0.163***	0.163***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Age^2	-0.002***	-0.002***	-0.002***	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Doctor		-0.084**	-0.166***	-0.101**	-0.115***
		(0.039)	(0.039)	(0.041)	(0.041)
Mgmt. Background			-0.105**	-0.093*	-0.106**
			(0.053)	(0.054)	(0.053)
Doctor × Mgmt. Studies				-0.199***	-0.199***
				(0.037)	(0.037)
Observations	8,197	8,197	8,197	8,197	8,185
R-squared	0.101	0.102	0.102	0.109	0.110
Sample	All	All	All	All	Degree data available

Notes: This table presents the correlation between the CEO fixed effects estimated from Equation B.1 and manager characteristics. These characteristics include gender, age, age^2, and a set of indicators for educational background. "Mgmt. Background" refers to undergraduate studies in management and "Mgmt. Studies" refers to postgraduate studies related to management. Controls include connected set fixed effects. Robust standard errors in parentheses.

Appendix C

Equilibrium Effects of Food Labeling Policies

C.1 Changes in prices, product assortment, and package size

In this appendix, we study how and whether firms responded to the policy by changing prices, product assortment, or package size.

To quantify the effects of the policy on equilibrium prices, we follow the event study strategy implemented for changes in equilibrium quantities from Equation (3.1). We estimate the following regression:

$$\log(p_{jst}) = \sum_k \beta_k \cdot L_j \cdot \mathbb{1}\{k = t\} + \delta_{js} + \delta_t + \varepsilon_{jst} \tag{B.1}$$

where all variables and specification details are defined as in Equation (3.1). Results are presented in Figure C.1.1, Panel (a). We find that labeled products saw an average decrease of 5.5% in prices relative to unlabeled products. This may be explained by a combination of firms increasing markups on unlabeled products that now face higher demand (and vice versa), and by an increase in marginal costs of unlabeled products due to reformulation. It might also be the case that firms are decreasing prices of labeled products due to their lower demand.

The previous result must be taken with caution, as prices could change due to a change in the mix of UPCs offered for a given product (e.g., changes in package sizes), and not because the offered price changes. In Figure C.1.1, Panel (b), we show the same coefficients from Equation (B.1) but aggregate the data at UPC-level. Using this specification, we find that labeled UPCs saw an average decrease of 4.2% in prices relative to unlabeled UPCs.

These results are in contrast to those in Pachali et al. (2022), who conclude that warning labels lead to higher prices of labeled cereals due to changes in product differentiation. The differences seem to be driven by differences in the sample. While we use scanner data from

Figure C.1.1: Event study for cereal prices

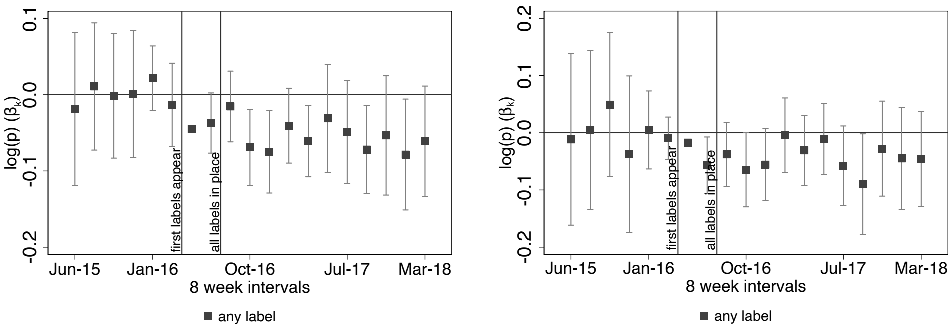

(a) Product-level data (b) UPC-level data

Notes: This figure presents the β_k coefficients of our event study regression for prices from Equation (B.1). Vertical segments delimit the 95% confidence intervals. Panel (a) uses product-level data and is estimated on the sample of 68 ready-to-eat cereals that show up in the pre- and post-policy periods. The sample consists of 27 unlabeled and 41 labeled products. Panel (b) uses UPC-level data and is estimated on a sample of 257 unique UPCs in the cereal market. The sample consists of 86 unlabeled and 135 labeled UPCs.

Walmart, they use household panel data from Kantar World-panel Chile. Moreover, of the 94 products in our sample, they focus on 14, of which only three are unlabeled. When repeating the analysis in our data but restricting it to the 14 products in their sample, we find no significant differences in price changes between labeled and unlabeled products.

We then study how product variety changed at Walmart before and after the policy implementation. We measure product variety by looking at the number of different products offered in each supermarket at a given period of time. To this end, we run the following regression:

$$\log(N_{st}) \quad = \quad \beta_t + \delta_s + \varepsilon_{st}, \tag{B.2}$$

where N_{st} is the total number of different products offered in store s in period t, and β_t and δ_s are period and store fixed effects, respectively. In Figure C.1.2, Panel (a), we plot the resulting coefficients β_t. We find that the number of products available increased by around 40% during the whole sample. Nevertheless, it does not seem that the increase in variety is directly related to the policy. No product was discontinued in our sample.

Finally, we look at changes in package size. Previous literature has suggested that policies that increase consumer attention to nutritional information can lead to reductions in package or serving size (Mohr, Lichtenstein, and Janiszewski, 2012). It is important to notice that in such settings, nutritional content is usually reported on a "per-serving-size" basis. In the context of Chile, the labeling status of products depends on the sugar and caloric concentration per 100 grams of cereal, thus eliminating the incentive to manipulate package or serving size. To study what happened to the average size of the package after the policy

Figure C.1.2: Changes in product assortment and package size

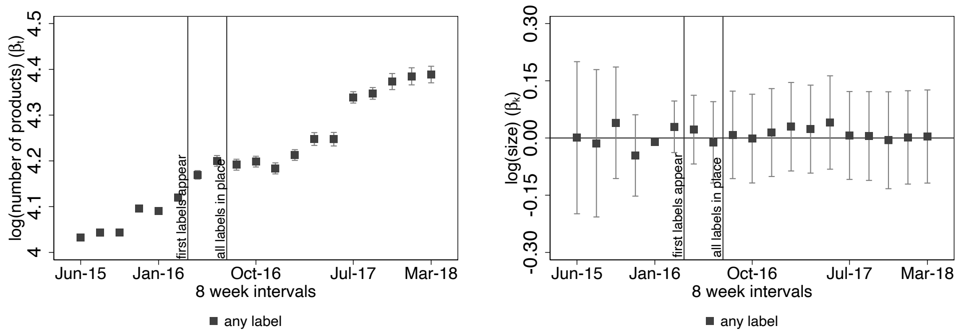

(a) Average number of products per store

(b) Average package size of products

Notes: This figure presents the β_t and β_k coefficients of the regressions from Equations (B.2) and (B.3). The vertical segments delimit the 95% confidence intervals. Panel (a) uses store-period-level data on a sample of 164 different stores. Panel (b) uses UPC-store-period-level data and a sample of 257 unique UPCs.

was implemented, we run the following regression:

$$\log(\text{size}_{ist}) = \sum_k \beta_k \cdot L_j \cdot \mathbb{1}\{k = t\} + \delta_{js} + \delta_t + \varepsilon_{jst} \qquad \text{(B.3)}$$

where size_{ist} is the size of the package for product j's UPC i in store s in period t. All other variables and details are defined as in Equation (3.1), and observations are at the UPC level. Results are presented in Figure C.1.2, Panel (b). We find that once the policy is implemented, there is no significant change in the average size of product packages.

C.2 Demand Model Discussion

Stockpiling

We assume static demand. However, cereal is a storable product, which can lead to dynamic incentives that can bias our estimates. Hendel and Nevo (2006a) show that ignoring such dynamics can lead to overestimates of own-price elasticities. We implement several tests for stockpiling behavior proposed by Hendel and Nevo (2006b). We find evidence in favor of stockpiling; however, the effects are much smaller than in Hendel and Nevo (2006b).

Throughout our analysis, we focus on within-consumer predictions and patterns of stockpiling behavior. We construct a dataset in which each observation is a cereal purchase made by a given household. For each observation, we calculate the number of days that passed since the last time the household purchased cereal as well as the number of days until the household's next cereal purchase. We also document whether the product purchased was on sale or not at the time of the purchase.

Assessing whether consumers stockpile in response to price movements would be straightforward if consumers' inventories were observed. For instance, we could test whether end-of-period inventories are higher after sales. However, consumption, and therefore inventories, are unobserved. Hendel and Nevo (2006b) propose a model of stockpiling with different implications that can be tested without requiring us to observe inventories. Specifically, we estimate the following model:

$$y_{it} = \beta sale_{it} + \delta_i + \epsilon_{it},$$

where $sale_{ijt}$ takes the value of one if household i purchases a cereal product in period t that was on sale. Coefficients δ_i control for household fixed effects. We test for the following implications under stockpiling behavior:

1. Duration until the following purchase is longer during a sale.
2. Duration from the previous purchase is shorter for purchases during a sale.
3. Non-sale purchases have a higher probability that the previous purchase was not during a sale.

To test for the first implication, we define the outcome variable as the number of days it took to household i to buy cereal again after their purchase in period t. Under stockpiling, we expect β to be *positive*. In Table C.2.1, Panel A, we find that $\beta = 0.877$, implying a 2.4% increase in the number of days until the next purchase when the product purchased is on sale. This number is positive but smaller in magnitude than those in Hendel and Nevo (2006b), who find a 10.6% and 9.3% increase in the market for yogurt and soft drinks, respectively.

To test for the second implication, we define the outcome variable as the number of days that passed since the last time household i purchased cereal before buying cereal again in period t. Under stockpiling, we expect β to be *negative*. In Table C.2.1, column (2), we find that $\beta = -0.420$, implying a 1.1% decrease in the number of days since the last purchase when the product purchased is on sale. This number is negative but smaller in magnitude than those in Hendel and Nevo (2006b), who find a 4.6% and 12.0% decrease in the market for yogurt and soft drinks, respectively.

To test for the third implication, we define the outcome variable to take the value 1 if household i's cereal purchase before buying cereal again in period t was of cereal products that were not on sale. Under stockpiling, we expect β to be *negative*. In Table C.2.1, column (3), we find that $\beta = -0.0633$, implying a 7.7% decrease in the probability that the last purchase was a non-sale purchase. This number is negative but smaller in magnitude than those in Hendel and Nevo (2006b), who find a 16.7% and 13.5% decrease in the market for yogurt and soft drinks, respectively.

Table C.2.1: Stockpiling tests

	(1) Days to next purchase	(2) Days since last purchase	(3) Prob of non-sale purchase
$sale_{it}$	0.877***	-0.420***	-0.0633***
	(0.041)	(0.041)	(0.0003)
Mean of dep. var.	37.00	37.00	0.81
Observations	10,580,676	10,580,676	10,580,676

Notes: In this table we present results of different test for stockpiling. In column (1), we test whether the duration until the following purchase is longer during a sale. In column (2), we test whether the duration from the previous purchase is shorter for purchases during a sale. In column (3), we test whether non-sale purchases have a higher probability that the previous purchase was not during a sale. We find evidence in favor of stockpiling; however, the effects are much smaller than found in other settings. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Our results are in line with O'Connell and Smith (2021), who perform similar tests in the soft-drinks market in the UK and find that the sign of these tests are consistent with stockpiling but very small in magnitude.

Salience effects

In this subsection, we investigate the potential salience effects of food labels in the cereal market. Salience refers to a situation in which an attribute of an item attracts more attention, and subsequently receives more weight when making decisions. In Section 3.3, we argue that labels shift consumer demand because they provide consumers with information about the true nutritional content of a product. However, labels may also make the unhealthiness of products more salient to consumers. In other words, labels may induce consumers to pay more attention to the role of sugar and calories in the decision-making process. Hence, if labels were only impacting demand through salience, we should expect the reduction in equilibrium quantities documented in Figure 3.2(a) to be stronger for those products with higher concentrations of critical nutrients.

To investigate this hypothesis, we follow the same empirical design implemented in Section 3.3. We split our sample of labeled products into two groups: products below the median in the caloric concentration distribution (20 products) and products above the median in the caloric concentration distribution (21 products). We use indicator dummies for each of these groups (denoted by Low_j^c and $High_j^c$) and estimate the following equation:

$$\log(q_{jst}) = \sum_k (\beta_k^l \cdot L_j \cdot Low_j^c + \beta_k^h \cdot L_j \cdot High_j^c) \cdot \mathbb{1}\{k = t\} + \gamma \cdot \log(p_{jst}) + \delta_{js} + \delta_t + \varepsilon_{jst}, \quad \text{(B.4)}$$

where all variables and specification details are defined as in Equation (3.1).

Figure C.2.1: Changes in equilibrium quantities by caloric concentration

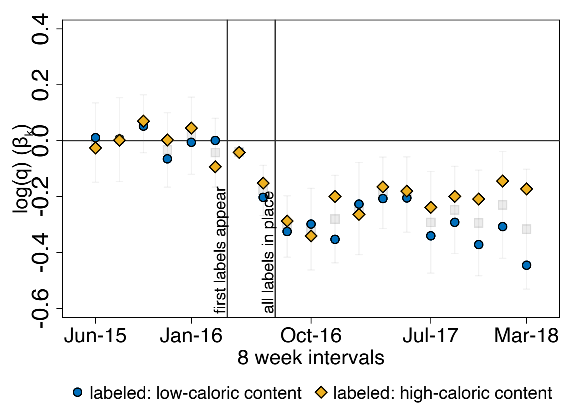

Notes: This figure displays the coefficients from Equation (B.4). Coefficients in blue circles and yellow diamonds denote β_k^l, β_k^h, respectively. Gray squares denote the β_k coefficients from Equation (3.1) and the vertical lines delimit their 95% confidence intervals. These regressions are run on the sample of 68 ready-to-eat cereals that show up in the pre- and post-periods. The sample contains 27 unlabeled products and 41 labeled ones.

Results from Equation (B.4) are shown in Figure C.2.1. Coefficients in blue and yellow denote β_k^l and β_k^h estimates, respectively. Coefficients in light gray denote β_k coefficients from Equation (3.1). Products with low caloric concentration (blue dots) and high caloric concentration (yellow diamonds) saw a similar reduction in equilibrium quantities.[1] If anything, high-calorie products seem to experience lower reductions in demand, as opposed to what we would expect under strong salience effects.

[1] Splitting products according to sugar concentration is less interesting. Because sugar concentration is highly correlated with beliefs about caloric concentration, results are similar to Figure 3.2(b). Labeled products with high sugar concentration experienced lower changes in equilibrium quantities than labeled products with low sugar concentration. This, again, rejects important salience effects.

Invariant taste

Equation (3.4) from the main article does not allow for the experience aspect of the utility, δ_{ijt}, to change when firms reformulate products and change w_{jt}. However, it could be the case that reducing the amount of calories or sugar in products renders them less appealing to consumers due to changes in taste.

In this subsection, we estimate a version of our demand model that allows for w_{jt} to directly affect the experience/taste aspect of consumers' utility function. Similar to the model in the main article, we assume that the utility derived by individual i when purchasing product j can be split into three main components:

$$u_{ijt} = \underbrace{\delta_{ijt}}_{\text{experience/taste}} - \underbrace{\alpha_i p_{jt}}_{\text{price paid}} - \underbrace{w'_{jt}\phi_b}_{\text{health consequences}} . \tag{B.5}$$

The main and most important difference between this model and the model in the main articule relies in the parameterization of the experience/taste aspect of the utility. In this section, we will allow δ_{ijt} to vary with w_{jt}. In particular, we assume that

$$\delta_{ijt} = w'_{jt}\gamma_b + \beta_i r_j + \delta_{jb} + \delta_{T(t)b} + \delta_{S(t)b} + \xi_{jtb} + \epsilon_{ijt}. \tag{B.6}$$

Consumers' decision utility in this model is then given by

$$\mathbb{E}_b[u_{ijt}] = -\alpha_b p_{jt} - \mathbb{E}_b[w_{jt}|L_{jt}]'\phi_b + w'_{jt}\gamma_b + \beta_i r_j + \delta_{jb} + \delta_{T(t)b} + \delta_{S(t)b} + \xi_{jtb} + \epsilon_{ijt}, \tag{B.7}$$

where ϕ_b determines changes in preferences driven by changes in beliefs about the nutritional content of a product and γ_b determines changes in preferences driven by the actual change in nutritional content of the product. Note that preferences driven by baseline beliefs and nutritional content are absorbed by product fixed effects δ_{jb}. Also note that consumers could respond to changes in w_{jt} even if w_{jt} is not observed by them but is correlated with things they do observe but the econometrician doesn't (e.g., taste). This model departs from the one estimated in Section 3.4 in two ways: First, we allow utility to directly depend on nutritional content w_{jt} through the term $w'_{jt}\gamma_b$. Second, we fix $\Sigma_\phi = \sigma_\alpha = 0$, which allows for more transparent identification of ϕ_b and γ_b. In a model in which consumers dislike a higher concentration of critical nutrients due to the negative health consequences of consuming them—but in which sugar and calories increase the taste of the products—we should expect to find that $\phi_b > 0$ and $\gamma_b > 0$.

There are two important challenges when trying to separately identify ϕ_b and γ_b. First, changes in nutritional content happen around the time of the policy implementation, and therefore changes in $\mathbb{E}_b[w_{jt}|L_{jt}]$ and w_{jt} happen at the same time. Second, changes in $\mathbb{E}_b[w_{jt}|L_{jt}]$ are not directly observed in the data. We infer them by combining the beliefs survey and a Bayesian updating model. If $\Delta\,\mathbb{E}_b[w_{jt}|L_{jt}]$ and Δw_{jt} are correlated and the former is measured with error, γ_b could capture parts of the effects driven by changes in beliefs.

In Figure C.2.2 we plot the changes in beliefs estimated in Section 3.4 of the main article vs. the changes in nutritional intake observed in the data for both sugar and calories. For both critical nutrients, there are products for which nutritional content changed but beliefs did not (products that were believed to be low in sugar or calories and that had to reformulate to avoid receiving the label) as well as products for which nutritional content did not change but beliefs did (products that were believed to be low in sugar or calories but did not reformulate and received a label). We exploit changes in demand for both types of products to separately identify ϕ_b and γ_b.

Figure C.2.2: Changes in beliefs vs. changes in real nutritional content

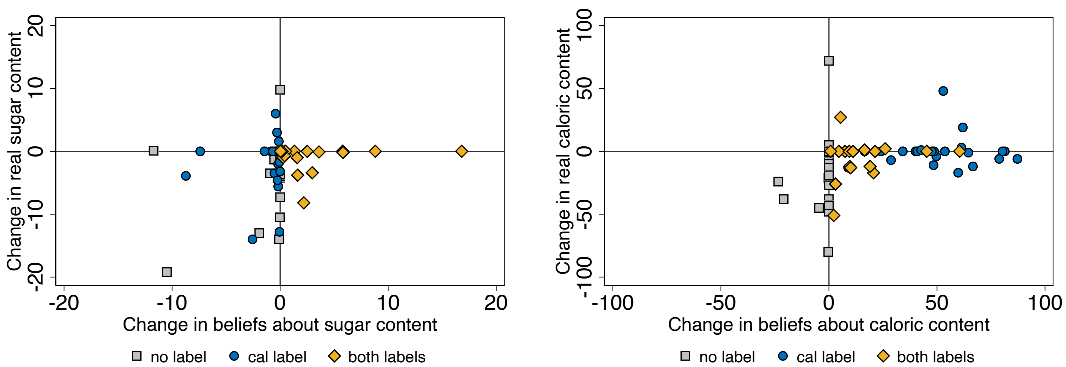

(a) Sugar

(b) Calories

Notes: The figure shows changes in beliefs about nutritional content vs. changes in real nutritional content. To calculate changes in beliefs about nutritional content, we subtract the estimates of $\mathbb{E}_b[w_{jt}|L_{jt}]$ from before and after the policy implementation. We calculate changes in real nutritional content directly from the data. Gray squares are products that did not receive any label, blue circles are products that received a high-in-calorie label, and yellow diamonds are products that received both a high-in-calorie and a high-in-sugar label. Panel (a) shows results for sugar and Panel (b) shows results for calories.

To estimate the model, we fix the nonlinear parameters μ, Σ_β, and ρ at the estimated values of the model from Section 3.4 in order to keep both models as close as possible. We also add additional instruments for the identification of γ_b by interacting the pre-policy nutritional content with dummies for whether a given product was above or below the threshold and with a dummy for the post-policy period. The intuition behind the instrument is that products above the threshold in the pre-policy period are more likely to reformulate, and products that bunch and are closer to it will reduce their nutritional content less than those that bunch but are further from it.

We present the results in Table C.2.2. The parameter estimates show that higher concentrations of sugar and calories do not imply higher taste, thus rejecting the hypothesis that reformulated products substantially decreased their taste. This is consistent with the evidence provided in Appendix C.3, in which we explain that the reformulation process took place with the explicit goal of not affecting the product's taste. More surprisingly, we find that $\gamma_b^c < 0$, which implies that reducing caloric content increases the taste of the product. We believe this finding is driven by measurement error in the change in beliefs shown in Figure C.2.2. Products that, on average, were believed to be low in calories and reformulated calories to avoid receiving the label should see no changes in beliefs, according to our model. However, some consumers may be learning from the labels regardless, which can induce increases in demand for those products despite reducing their calories.

Table C.2.2: Estimated demand parameters with variable taste

Panel A: Preferences for price and healthiness (α_b)				
	low-SES		high-SES	
Price (α_b)	α_l	0.2759***	$\bar{\alpha}_h$	0.2086***
		(0.0200)		(0.0221)

Panel B: Preferences for healthiness and taste (ϕ_b, γ_b)								
	Sugar				Calories			
	low-SES		high-SES		low-SES		high-SES	
Healthiness (ϕ_b)	ϕ_l^s	0.0054**	ϕ_h^s	0.0045	ϕ_l^c	0.0387***	ϕ_h^c	0.0369***
		(0.0028)		(0.0031)		(0.0034)		(0.0042)
Taste (γ_b)	γ_l^s	-0.0033	γ_h^s	0.0010	γ_l^c	-0.0176***	γ_h^c	-0.0221***
		(0.0029)		(0.0036)		(0.0057)		(0.0071)

Notes: This table shows the main results from estimating the model from Equation (B.7). Nutritional content is measured in grams of sugar and kilocalories per gram of cereal and prices in dollars per 100 grams of cereal. Standard errors are reported in parentheses.

Advertising

Our model does not account for potential changes in advertising due to the labeling policy. The Chilean Food Act imposed additional marketing restrictions by not allowing firms to advertise labeled products to children under age 14 across different platforms, including websites, social media, magazines, billboards, pamphlets, newspapers, radio, and television. Correa, Reyes, Taillie, Corvalán, and Dillman Carpentier (2020) show that the policy was effective in decreasing advertising of labeled products by documenting a decrease in the share of food advertising that includes labeled products from 41.9% of total food advertising in the pre-policy period to 14.8% in the post-policy period. Since changes in advertising are potentially correlated with changes in beliefs, some of the effects we attribute to changes in beliefs may be driven by changes in advertising. In this subsection, we use data collected by Correa et al. (2020) and show that all of our estimates are robust to including TV advertising intensity in the utility function.

The data we use comprise all television ads aired on the four main broadcast channels in Chile during a stratified random sample of days in April and May of 2016 (pre-policy) and 2017 (post-policy). Of all ads during the pre-policy period, only 0.5% displayed a product belonging to the breakfast cereal category. Moreover, 9 products appeared in an ad in the pre-policy period and only 6 in the post-policy period. The average number of ads per product on a given day and channel, once we condition for those products that appeared in any ad, is 0.3. This already suggests that the role of TV advertising in the cereal market is likely to be small.

To empirically test whether advertising bans played an important role in consumer choices, we add an additional element to consumers' decision utility:

$$\mathbb{E}_b[u_{ijt}] = -\alpha_i p_{jt} - \mathbb{E}_b[w_{jt}|L_{jt}]'\phi_i + \gamma_b A_{jt} + \beta_i r_j + \delta_{jb} + \delta_{T(t)b} + \delta_{S(t)b} + \xi_{jtb} + \epsilon_{ijt}, \quad \text{(B.8)}$$

where A_{jt} is a measure of advertising intensity for product j in market t, and all other variables are the same as in the model from Section 3.4 in the main article.[2] We measure advertising intensity as the average daily number of ads shown on each TV channel for each product.[3] Since we only have two snapshots of advertising intensity, we follow the same strategy used for reformulation and changes in beliefs, and assume that all changes happened at the time of the policy implementation. We present the results in Table C.2.3.

All coefficient magnitudes are almost identical to the main specification in the text. Moreover, the coefficients on γ_b are small in magnitude and not statistically different from zero. Our estimates imply that consumers are willing to pay between \$0.032 and \$0.044 more per 100 grams of cereal for each additional ad shown on every channel, every day.

[2] We estimate the model following the same methodology as in Section 3.4, including A_{jt} interacted with consumer type dummies as additional instruments.

[3] Our results are robust to other measures of advertising, such as average daily ad minutes per channel and average daily minutes times rating points per channel.

Table C.2.3: Estimated demand parameters with advertising

Panel A: Preferences for price and healthiness (α_i, ϕ_i)								
	First moments				Second moments			
	low-SES		high-SES		low-SES		high-SES	
Price (α_i)	$\bar{\alpha}_l$	0.2517***	$\bar{\alpha}_h$	0.1864***	σ_{α_l}	0.1504***	σ_{α_h}	0.1114***
		(0.0733)		(0.0597)		(0.0337)		(0.0359)
Sugar (ϕ_i^s)	$\bar{\phi}_l^s$	0.0129***	$\bar{\phi}_h^s$	0.0129**	$\sigma_{\phi_l^s}$	0.0414	$\sigma_{\phi_h^s}$	0.0415
		(0.0043)		(0.0052)		(0.1115)		(0.1120)
Calories (ϕ_i^c)	$\bar{\phi}_l^c$	0.0261***	$\bar{\phi}_h^c$	0.0254***	$\sigma_{\phi_l^c}$	0.0278	$\sigma_{\phi_h^c}$	0.0271
		(0.0075)		(0.0078)		(0.0181)		(0.0171)

Panel B: Individual preferences for different subcategories (Σ_β)									
Plain		Sugary		Chocolate		Granola		Oatmeal	
$\sigma_{\beta_{r_1}}$	0.0577	$\sigma_{\beta_{r_2}}$	0.1991	$\sigma_{\beta_{r_3}}$	0.2077	$\sigma_{\beta_{r_4}}$	0.0350	$\sigma_{\beta_{r_5}}$	0.2828
	(0.1463)		(0.1887)		(0.1355)		(0.1633)		(0.3513)

Panel C: Nest, beliefs, and advertising parameters (ρ, μ, γ_b)							
				low-SES		high-SES	
Nest parameter	ρ	0.9607***	Advertising (γ_b)	γ_l	0.00810	γ_h	0.00813
		(0.0040)			(0.00706)		(0.00807)
Beliefs shifter	μ	-0.1255***					
		(0.0191)					

Notes: This table shows the main results from estimating the model from Equation (B.8). Nutritional content is measured in grams of sugar and kilocalories per gram of cereal and prices in dollars per 100 grams of cereal. Advertising intensity is measured as the average daily number of ads per channel for each product. For random parameters $x_i \in \{\alpha_i, \phi_i, \beta_i\}$, we report their average $\bar{x}$ and standard deviation σ_x. Standard errors are calculated using the delta method and reported in parentheses.

C.3 Supply Model Discussion

Timing of firms' choices

In the main article, we assume that firms choose prices and nutritional content simultaneously. In practice, firms are likely to first set the nutritional content of their products in their production facility and then choose prices in the retail stores. Due to strategic incentives, firms may want to deviate from $w_j = \nu_j$ even in the absence of regulation to increase the marginal cost and promote overall higher prices. Whether this incentive exists depends on the specific parameters and shape of the demand function. Here, we show that under a simple oligopolistic model with Bertrand competition, single-product firms, and logit demands, such incentive never arises. Then, we use simulations to show that in the more complicated setting of our framework with random coefficients and multi-product firms, no firm also has an incentive to deviate from $w_j = \nu_j$ in the absence of regulation.

First, note that in our model, demand $s_{jt}(\mathbf{p}, \mathbb{E}_\pi[\mathbf{w}|\mathbf{L}])$ does not directly depend on w_j in the absence of regulation. Therefore, the problem of choosing nutritional content w_j is equivalent to the question of setting marginal cost c_j when marginal cost does not enter in the demand function. In the simultaneous game, it is straightforward to show that, from the first-order conditions, firms set costs at the minimum possible value (see Section 3.5 of the main article). We show next that in a sequential model with single-product firms and logit demand, in which firms set marginal cost first and then choose prices, it is also an equilibrium for all firms to choose the minimum cost.

Let the profit function of a single-product firm be given by $\pi_j(\mathbf{p}, c_j) = (p_j - c_j)s_j(\mathbf{p})$, where $s_j(\mathbf{p}) = \frac{exp(-\alpha p_j + \delta_j)}{1 + \sum_k exp(-\alpha p_k + \delta_k)}$. In the first stage of the sequential model, firms choose $c_j \geq \underline{c}_j$. In the second stage, after marginal costs are realized, firms choose p_j.

First, note that under logit demand, $\pi_j(\mathbf{p}, c_j)$ has increasing differences in (p_j, p_{-j}), which means that the second-stage game in the sequential model is a supermodular game. Also, note that $\pi_j(\mathbf{p}, c_j)$ has increasing differences in (p_j, c_j), which implies that larger choices of c_j in the first stage will translate into larger choices of p_j in the second stage.

Let $\mathbf{p}^*$ be the vector of equilibrium prices in the second stage when all firms play $c_j = \underline{c}_j$ in the first stage. We want to show that no firm j has incentives to deviate and choose $c_j > \underline{c}_j$ in the first stage.

Suppose that j deviates and chooses $c'_j > \underline{c}_j$ in the first stage. Let p'_j be the price specified by j's strategy following such a deviation, and $\mathbf{p}'$ the equilibrium price vector after the deviation. Because $\pi_j(\mathbf{p}, c_j)$ has increasing differences in (p_j, c_j), we know that $p'_j \geq p^*_j$. Moreover, because the second-stage game in the sequential model is a supermodular game, we will also have that $\mathbf{p}' \geq \mathbf{p}^*$ (i.e., all firms will set larger prices in the second stage after the deviation).

From the first-order conditions of firm k, we have that $s_k(\mathbf{p}') \geq s_k(\mathbf{p}^*)$. It is also straightforward from the logit demand formula that $s_0(\mathbf{p}') \geq s_0(\mathbf{p}^*)$, where $s_0(\cdot)$ is the market share of the outside option. Because market shares add up to one, we have then that $s_j(\mathbf{p}') \leq s_j(\mathbf{p}^*)$. Finally, with logit demand, lower market shares imply lower markups.

Thus, we have that $\pi_j(\mathbf{p}', c'_j) \leq \pi_j(\mathbf{p}^*, \underline{c}_j)$, which proves that firm j has no incentive to deviate.

We test this result in the context of our estimates using the simulations from the counterfactual analysis of Section 3.6. For each simulation, we ask each firm whether they would be willing to deviate from $w_j = \nu_j$ in a potential first stage. We find that no firm would increase their profits by implementing such deviation.

Comparing the simultaneous and sequential games when a labeling policy is in place is more complicated due to the potential presence of multiple equilibria. In our simulations, we find that whether a firm decides to bunch or not is mostly driven by Λ_j, the cost of decreasing a product's nutritional content. Products with a low value of Λ_j tend to always reformulate, while products with a large value of Λ_j never reformulate. Because the decision to bunch is discrete, a firm's optimal response is constant under a large range of strategies p_{-j}. This means that in our setting, the equilibrium tends to be unique and identical in both the simultaneous and sequential games.

Reformulation process

In the main article, we assume that reformulation does not change the taste of products. This assumption simplifies the firm's problem of choosing w_{jt} in the absence of regulation, which we use to estimate ν_j from the first-order conditions. This assumption is driven by industry participants' descriptions of how reformulation was accomplished which we describe below. We also assume that reformulation changes marginal cost and do not model it as a fixed cost. This is consistent with how reformulation operated in the cereal market, where the techniques used were already developed in other countries and widely used in the diabetic food industry.

There are two potential ways firms may reformulate their products. In one way, firms may choose to sacrifice taste for healthiness by removing some of the critical nutrients from their products. In the other way, firms may choose to replace critical nutrients with alternative, potentially more expensive, ingredients without compromising taste, mouthfeel, shelf life, and other attributes to ensure that consumers will continue to buy their products.

We conducted interviews with consumer product managers at the two largest ready-to-eat cereal producers in Chile and asked them about their reformulation process. They explained that when products are reformulated, it is an explicit goal of the company to produce products that are indistinguishable from the previous version. When making modifications to products, they follow different steps to ensure their goals are met. First, they hire a group of "taste experts" who work closely with the firm during the reformulation process and check that attribute standards are met. Then, they implement randomized blind tests to corroborate that consumers cannot distinguish between the old and new versions of the product. Only if a product successfully passes the different tests will firms release the new version of the product to the market.

Reformulating cereal products presents different challenges. One of the main roles of sugar is to deliver sweetness. Artificial and natural high-intensity sweeteners are alternatives to sugars (e.g., sucralose acesulfame-K, saccharin steviol glycosides). Firms usually also use taste enhancers to amplify the sweetness intensity of sweeteners like sucralose or stevia. Another key role of sugar in the production process is to provide volume and structure to cereals which artificial sweeteners do not. Without sugar, cereals crumble. Polyols, which are widely used in the diabetic food industry, act as bulking agents and provide thickness and structure to products. They are less sweet than sucrose and deliver a clean, non-lingering sweet taste very close to the profile of sucrose. Combinations of polyols with intense sweeteners and/or sweetness enhancers allow a higher level of sweetness intensity while maintaining the important physicochemical properties of sugars (Lê, Robin, and Roger, 2016). Replacing sugar with these ingredients results in a more expensive product to produce, which raises the cost of cereal ingredients by more than 20%, according to the product managers.

We collected data on the specific ingredients of 17 of the 20 products that were reformulated in our sample. We found that after the policy is implemented, 47% start using maltitol (a type of polyols), 29% sucralose, and 35% stevia.

C.4 Additional Figures and Tables

Figure C.4.1: Distribution of caloric and sugar concentration pre- and post-legislation

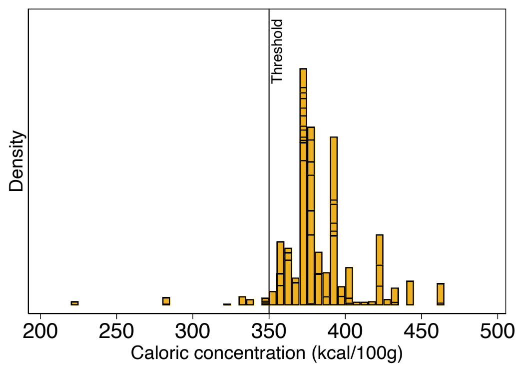

(a) Distribution of calorie content in 2016

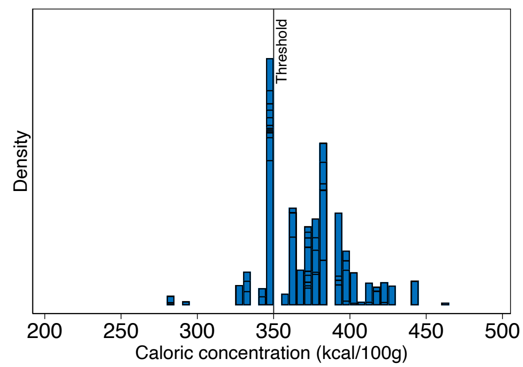

(b) Distribution of calorie content in 2018

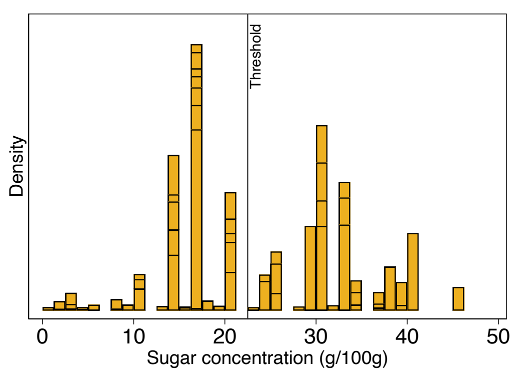

(c) Distribution of sugar content in 2016

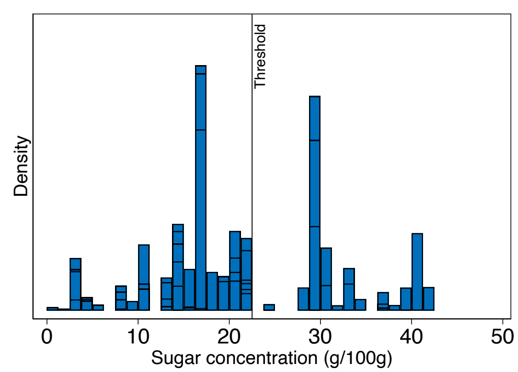

(d) Distribution of sugar content in 2018

Notes: This figure plots the distribution of calories and sugar per 100g for cereal products before and after the policy implementation. Horizontal black lines inside the bars identify different products. Observations are weighted by pre-policy revenue. We exclude oatmeal products, which do not have artificially added critical nutrients, as they are exempted from the regulation and do not reformulate their products.

Figure C.4.2: Distribution of markups

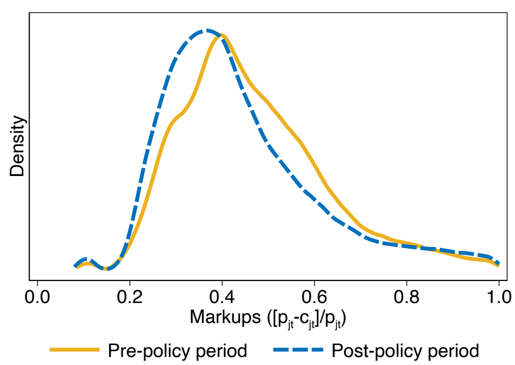

Notes: This figure shows the distribution of markups—defined as the ratio of price minus marginal cost to price—across products and markets before and after the policy implementation.

Figure C.4.3: Changes in consumer welfare under different values of ρ

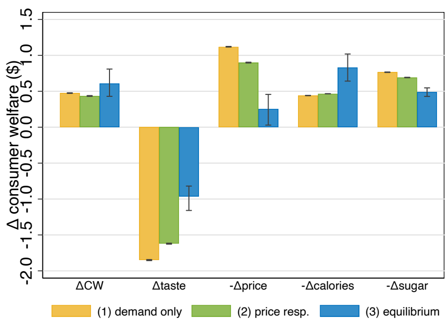

(a) $\rho = 0.9$

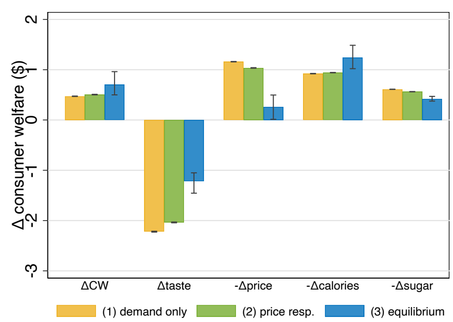

(b) $\rho = 0.8$

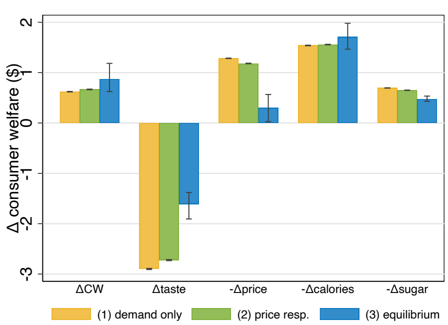

(c) $\rho = 0.7$

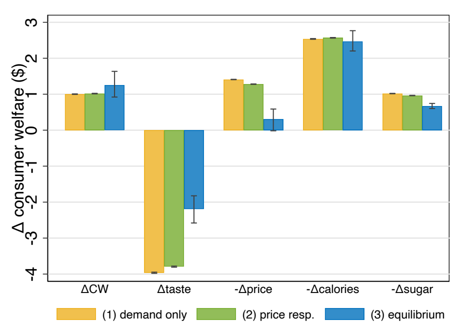

(d) $\rho = 0.6$

Notes: This figure replicates the findings from Figure 3.6 by imposing different values of ρ. For each panel, we fix ρ at 0.9, 0.8, 0.7, and 0.6, respectively. For each value of ρ, we estimate all other parameters from the demand and supply models presented in Sections 3.4 and 3.5. We then run our main counterfactuals and calculate the changes in consumer welfare under the different parameters. We show that our main results are qualitatively similar when we assume lower values of ρ.

Figure C.4.4: Beliefs about nutritional content vs true post-policy nutritional content

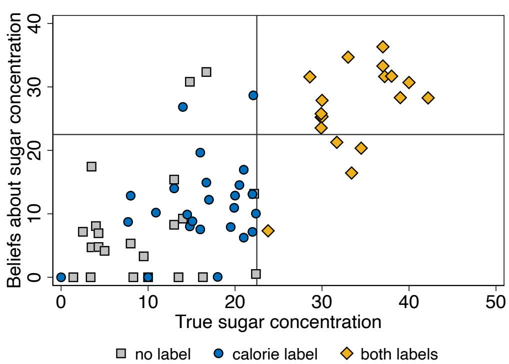

(a) Sugar

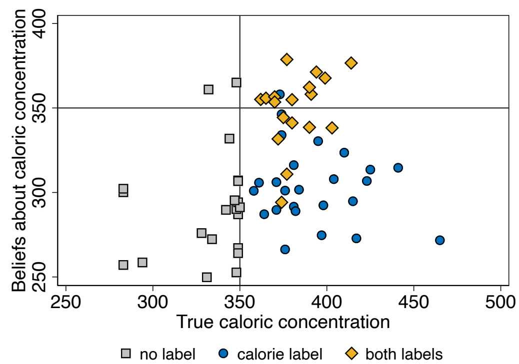

(b) Calories

Notes: This figure shows the estimated average belief (between low- and high-SES consumers) about each product's nutritional content against the true post-policy period nutritional content. Vertical and horizontal lines correspond to the value of the policy threshold in both spaces. Gray-square products did not receive any label, blue-circle products received a high-in-calorie label, and yellow-diamond products received a high-in-calorie and a high-in-sugar label. We exclude products that do not show up in the pre-policy period or are exempt from the policy.

Figure C.4.5: Predicted probability of bunching as a function of prior beliefs

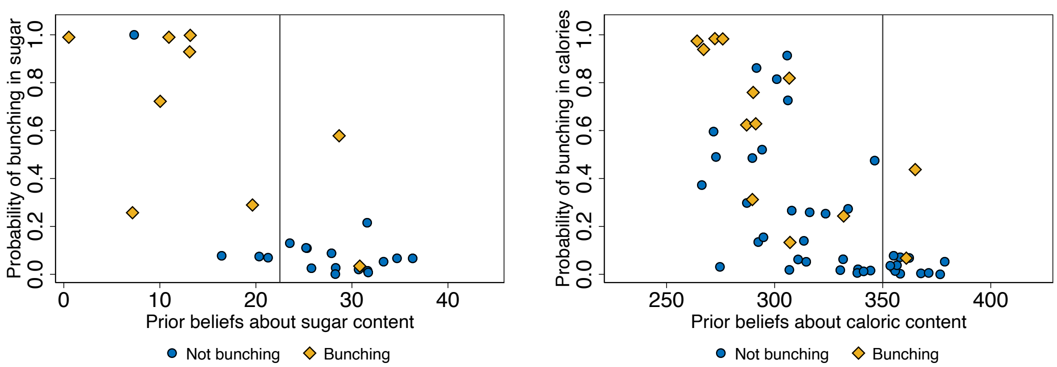

(a) Sugar

(b) Calories

Notes: The figure shows the predicted probability of each product bunching in sugar and calories as a function of the average prior belief about their nutritional content. In Panel (a), we focus on sugar content. Products in yellow diamonds are products that bunched in the data and crossed the sugar policy threshold. Products in blue circles are products that did not bunch and received a "high-in-sugar" label. In Panel (b), we focus on caloric content. Products in yellow diamonds are products that bunched in the data and crossed the calorie policy threshold. Products in blue circles are products that did not bunch and received a "high-in-calorie" label.

Table C.4.1: Median price elasticities

		Share (%)	Elasticities									
			Plain		Sugary		Chocolate		Oatmeal		Granola	
			(1)	(2)	(4)	(5)	(7)	(8)	(10)	(11)	(13)	(14)
Subcategory: Plain												
Fitness, Nestlé	(1)	0.77	-3.104	0.157	0.078	0.054	0.010	0.047	0.213	0.062	0.024	0.017
Quadritos, Quaker	(2)	0.66	0.173	-4.069	0.079	0.056	0.105	0.050	0.205	0.059	0.022	0.016
Corn Flakes, Nestlé	(3)	0.61	0.173	0.166	0.082	0.056	0.103	0.048	0.201	0.061	0.020	0.015
Subcategory: Sugary												
Trix, Nestlé	(4)	1.57	0.032	0.030	-3.224	0.484	0.014	0.007	0.068	0.020	0.005	0.004
Zucaritas, Kellog's	(5)	1.27	0.032	0.031	0.687	-3.113	0.015	0.007	0.063	0.018	0.005	0.004
Zucosos, Nestlé	(6)	0.69	0.032	0.030	0.673	0.486	0.015	0.007	0.068	0.020	0.005	0.004
Subcategory: Chocolate												
Chocapic, Nestlé	(7)	4.27	0.022	0.021	0.008	0.005	-1.996	0.361	0.066	0.019	0.003	0.003
Milo, Nestlé	(8)	1.55	0.022	0.021	0.008	0.005	0.779	-3.070	0.065	0.019	0.003	0.003
Mono Balls, Costa	(9)	0.94	0.020	0.019	0.007	0.005	0.778	0.356	0.064	0.018	0.003	0.002
Subcategory: Oatmeal												
Avena Instantanea, Quaker	(10)	5.8	0.084	0.073	0.066	0.043	0.121	0.055	-0.813	0.075	0.041	0.033
Avena Instantanea, Vivo	(11)	1.98	0.083	0.072	0.066	0.043	0.122	0.055	0.255	-0.831	0.041	0.034
Avena Tradicional, Quaker	(12)	1.55	0.084	0.073	0.066	0.043	0.122	0.055	0.255	0.074	0.041	0.033
Subcategory: Granola												
Granola Miel y Alm., Quaker	(13)	0.55	0.032	0.028	0.018	0.013	0.020	0.009	0.143	0.038	-2.880	0.223
Granola Miel y Alm., Vivo	(14)	0.45	0.029	0.023	0.018	0.012	0.020	0.009	0.145	0.042	0.291	-2.831
Granola Berries, Vivo	(15)	0.36	0.030	0.025	0.017	0.012	0.019	0.009	0.145	0.042	0.292	0.246
Outside option	(16)	61.08	0.001	0.001	0.002	0.001	0.003	0.001	0.002	0.001	0.001	0.000

Notes: The first column reports the median market share of each product across all 2,704 markets. For the rest of the table, cell entries j, k—where j indexes rows and k columns—give the percent change in market share of product j with a 1% increase in the price of product k. Each entry represents the median of the elasticities from all markets. Note that the cross-price elasticities within a subcategory are relatively constant. We do not observe product characteristics that vary within subcategories which limits our ability to include preference heterogeneity to recover more flexible substitution patterns.