Equilibrium Transition and Cartel Formation: A Structural Analysis of Chile's Pharmacy Cartel

Yu (Jasmine) Hao*

Faculty of Business and Economics

University of Hong Kong

haoyu@hku.hk

September 15, 2026

Abstract

This paper studies how Chile's three largest pharmacy chains moved from the price war to collusion, using court-record daily prices and a structural model. A court-ordered advertisement ban ended the comparison campaign, reducing the gain from loss-leader pricing and weakening the incentive to continue the war. With an upstream supplier's help, the chains **verified** a **procedure** for raising prices in turn. Once verified, they restored margins on former loss leaders and raised prices to extract rents, spreading more slowly than margin restoration. A model letting leaders' beliefs about follower participation update matches this pattern better than one imposing rational beliefs.

Keywords: cartel formation; dynamic games; adaptive learning; subjective beliefs; procedure verification; loss-leader pricing; advertisement ban; retail pharmacy.

JEL classification: L41, L13, M37, L81, K21, C57.

1 Introduction

How do firms start colluding? Although collusion has long been a central topic in industrial organization, most studies focus on its implementation and maintenance, with less attention

*I thank the Chilean Competition Authority (FNE) for data access; the editor and anonymous referees for comments; and my advisor, Hiro Kasahara, my dissertation committee, and Victor Aguirregabiria for discussions. I also thank Jean-François Houde for encouraging this study of the advertisement ban. I acknowledge support from the 2020 Gambling Award Fellowship. All errors are my own.

to cartel formation. Harrington (2018) distinguishes the problem of coordinating on a collusive agreement from that of ensuring compliance with it. Firms must establish a common understanding of what each will do and believe that their rivals will act accordingly. Byrne and De Roos (2019) provide one of the few empirical studies of this process, documenting how firms gradually establish a focal pricing cycle through price leadership and experimentation. I use a structural model to distinguish knowing how to coordinate a price increase from willingness to participate in subsequent increases. Verification makes the procedure common knowledge: one chain raises first, and the others decide in turn whether to follow. Knowing the procedure does not remove uncertainty about their participation. My analysis is also related to Igami and Sugaya (2022), who use a structural model to measure firms' incentives to sustain collusion and explain cartel breakdown. Here, the focus is on the build-up of collusion.

I study Chile's three largest pharmacy chains: Cruz Verde, Farmacias Ahumada (FASA), and Salcobrand. Together, they accounted for approximately 92 percent of national retail pharmacy sales in 2007.[1] I use the case records together with daily prices and quantities, wholesale costs, and manufacturers' suggested retail prices to study this transition. In late 2006, the chains began undercutting one another. They used loss-leader drugs to attract customers who would buy the discounted medicine and other medicines or products during the same pharmacy visit (TDLC, 2012, consids. 42–43). In August 2007, Cruz Verde launched a campaign comparing its prices with FASA's. Its rivals responded with further price cuts (FNE, 2008).

A court-ordered advertisement ban ended the comparison campaign on 6 November 2007. The sustained price war ended at the same time.[2] On 9 December 2007, with help from upstream suppliers, the chains first successfully used a "one-two-three" price-increase procedure. One chain raised its price first, and the other two followed in turn. The suppliers passed the proposed prices and dates between the chains (FNE, 2008; TDLC, 2012). The chains quickly used the same procedure for other medicines, raising prices toward the manufacturers' suggested retail levels and restoring positive margins. Prices rose from below

[1]The cartel case is *Fiscalía Nacional Económica v. Farmacias Ahumada S.A. and Others*, Rol C-184-08, before Chile's Competition Tribunal (TDLC). Judgment No. 119/2012 sanctioned Cruz Verde and Salcobrand; FASA had previously settled with Chile's competition authority, the Fiscalía Nacional Económica (FNE) (TDLC, 2012). The FNE helped me access the public case materials through Chile's transparency procedure.

[2]Santiago's 17th Civil Court ordered the campaign to stop in *Farmacias Ahumada S.A. v. Farmacias Cruz Verde S.A.*, case Rol C-23.423-2007 (FNE, 2008; 17º Juzgado Civil de Santiago, 2007).

wholesale cost to approximately 22–25 percent above it. **Further increases, which I call rent extraction,** raised prices to approximately 43–48 percent above wholesale cost, but spread more slowly. The competition authority began requesting information in late March 2008.

Before the first successful coordinated increase, the chains repeatedly tried and failed to raise prices together. Between 6 November and 8 December 2007, the price records show 42 failed medicine-level attempts, 39 started by Salcobrand. Salcobrand also asked pharmaceutical laboratories to help the chains coordinate increases (TDLC, 2012, consids. 95, 97). Once the chains succeeded on 9 December, they quickly repeated the procedure across other medicines to restore positive margins. In a 19 December email, Salcobrand described the completed sequence as a "procedure" and said it expected to repeat it with more products and laboratories (TDLC, 2012, consid. 95).

The chains then used the same procedure to raise prices further, moving beyond margin restoration to rent extraction. These increases spread more slowly, even though the chains had already used the procedure successfully. Over matched 65-day windows, the rate of first increases per eligible medicine-week was 1.83 times as high for margin restoration as for rent extraction (Section 3.4). Why did further increases spread more slowly once the procedure was established? A chain raising first still risked losing customers if its rivals did not follow.

I model the advertisement ban as changing two primitives: demand and the set of pricing procedures available to firms. I estimate a logit demand model that allows a chain to attract extra demand by offering a low price. The estimated response to ordinary price differences changes little after the advertisement ban, but the extra gain from offering a low price falls sharply. The low-price group consists of chains pricing a medicine sufficiently below the highest-priced chain. More than one chain can belong to this group. The outside good is not purchasing the medicine from any of the three chains. Holding current prices and recent price changes constant, membership in the low-price group raises purchase odds relative to the outside good by 23.0 percent before the ban and 6.4 percent afterward. Demand also responds to recent price changes. A chain loses more customers after raising its price if its rivals have not followed, creating a cost of moving first. Recent cuts have no clear additional effect after controlling for current prices and low-price-group membership. I use these demand estimates to calculate the gains and losses from changing prices, including complementary basket profit from other purchases associated with the additional medicine demand.

I develop a dynamic pricing model with three blocks: ordinary price competition, the advertisement campaign, and collusive price leadership. The advertisement campaign allows larger cuts and changes how often medicines are reviewed for a price adjustment. After the ban, laboratories provide opportunities to coordinate price increases. Before verification, Salcobrand knows the procedure while the other chains evaluate increases under ordinary price competition. After verification, one chain raises first and the other two decide in turn whether to follow. The leader's risk is that rivals decline and customers leave. Knowing the procedure does not tell the leader whether its rivals will join a further increase.

I compare two models of leader confidence in follower participation. Adaptive Confidence updates the weight the leader places on the expected profit from sequential follower choices relative to raising alone after each rent-extraction attempt. Unit Confidence imposes rational beliefs by fixing that weight at one. Followers choose whether to join in both models. I estimate four economic parameters in each model and two additional learning parameters in Adaptive Confidence by simulated method of moments, matching nine groups of moments to the data (Goettler and Gordon, 2011; Yang, 2020). By the end of the estimation period, Adaptive Confidence predicts that about 68 medicines reach the rent-extraction level, compared with 70 in the data and about 106 under Unit Confidence.

This article is related to several strands of literature. First, it contributes to the literature on cartel formation. Much of the collusion literature studies whether firms can sustain a profitable agreement (Bain, 1959; Green and Porter, 1984). Forming a cartel also requires firms to coordinate on what to do (Green et al., 2014; Harrington, 2018). Byrne and De Roos (2019); Byrne et al. (2026) document how firms move toward collusion through price leadership, experimentation, and bargaining through prices. Related studies examine how communication, pricing rules, and intermediaries help firms coordinate (Genesove and Mullin, 1999, 2001; Harrington, 2017; Gerlach and Nguyen, 2021; Miller et al., 2021). Chaves and Duarte (2025) study how fuel distributors helped coordinate price increases among gas stations, while Alé-Chilet and Atal (2020) show how a trade association raised physicians' prices by coordinating their exit from insurer networks. Studies of the Chilean pharmacy cartel document the initial use of safer, more differentiated products, consistent with mistrust (Alé-Chilet, 2016), and examine the costs and allocation of price leadership (Alé-Chilet, 2018). I build on this work by separating verification of a price-increase procedure from uncertainty about rivals' willingness to participate. The same procedure spread quickly during margin restoration but more slowly during rent extraction. The structural approach is also related to Igami and Sugaya (2022), who study incentives to sustain collusion under a known quota agreement. My focus is on the build-up of collusion, when firms first verify the procedure and then consider further increases while participation remains uncertain.

Second, this paper relates to the literature on loss-leader pricing. Loss-leader pricing uses below-cost products to attract customers who also buy other goods (Lal and Matutes, 1994; DeGraba, 2003; Chen and Rey, 2012). Profits from these cross-category purchases give the chains a reason to keep cutting medicine prices despite negative margins on the discounted medicines (Thomassen et al., 2017). I connect this mechanism to the literature on advertising restrictions (Benham, 1972; Cady, 1976; Milyo and Waldfogel, 1999). I use the demand estimates to assess how the ban weakened the incentive to continue the price war. The ban therefore reduced the gains from using low-priced medicines to attract customers, weakening the incentive to undercut.

Third, this paper relates to the literature on firm learning. The firm-learning literature studies learning about market conditions and competitors' behavior (Aguirregabiria and Jeon, 2020). Doraszelski et al. (2018) study how firms learn in a new electricity market using models of fictitious play and adaptive learning. Related work examines how retailers learn to set prices after liquor privatization (Huang et al., 2022) and how learning about demand affects investment in container shipping (Jeon, 2022). Aguirregabiria and Magesan (2020) allow firms' beliefs about rivals' actions to differ from actual behavior in a dynamic game of retail-chain location. In my model, verification makes the procedure common knowledge, but uncertainty about participation remains. Under Adaptive Confidence, the leader uses a Beta-shaped outcome-updating rule for its subjective weight after each rent-extraction attempt. Complete participation raises the weight, while incomplete participation lowers it.[3] Firms can therefore understand the same procedure yet remain uncertain about whether others will participate in further increases.

Section 2 presents the institutional setting and data. Section 3 documents the transition from the price war to collusion. Section 4 estimates demand and the return to undercutting. Section 5 presents and estimates the dynamic pricing model, separating procedure verification from beliefs about rival participation. Section 6 concludes.

2 Institutional Setting and Data

2.1 Market overview

The Chilean pharmacy industry was dominated by three chains. In 2007, Cruz Verde, the largest chain,[4] accounted for 40.6% of nationwide pharmacy sales, followed by FASA

[3] The reduced-form rule treats each selected rent-extraction attempt as a binary outcome: $y = 1$ for complete participation and $y = 0$ for incomplete participation. It starts from the pseudo-counts $\nu^r m_0^r$ and $\nu^r(1 - m_0^r)$, so after S_t^r complete and F_t^r incomplete outcomes it produces the weight in Equation (8).

[4] Cruz Verde was vertically integrated with the distributor Socofar (FNE, 2008); see Online Appendix A, "(b) Three chains setting one national price schedule."

with 27.7% and Salcobrand with 23.8%.[5] A contemporaneous industry report counted 494 Cruz Verde locations, including 153 franchises, 359 FASA locations, and 314 Salcobrand locations (Durán and Kremerman, 2007, pp. 21, 41, and 55). The FNE reported that most of Chile's roughly 600 independent pharmacies did not carry branded prescription medicines and therefore provided little competition to the three chains for those products (FNE, 2008, para. 82). Each chain centrally set a national price schedule, so major price increases and cuts were implemented chain-wide rather than store by store. All three chains sold many products besides medicines. A low medicine price could therefore attract shoppers who also bought other items.

According to the FNE, consumers could fill a prescription at any chain but needed medical supervision to switch medicines. The physician therefore largely determined the medicine, while the consumer chose the pharmacy. The FNE describes these consumers as "captive through the prescription" (*cautivos a través de la receta médica*) (FNE, 2008, paras. 42–44 and 74).[6] A 2006 report by Chile's consumer-protection agency (SERNAC) says that proximity was the main factor in pharmacy choice, followed by price, loyalty to a chain, and comparing prices before buying (Servicio Nacional del Consumidor, 2006). Whether covered by public or private health insurance, households still paid directly for retail medicines (Castillo-Laborde and Villalobos Dintrans, 2013). Medicines were the largest component of household out-of-pocket health spending (Debrott, 2007). I therefore model consumers' choice among chains for a given medicine.

Chile had separate public and private procurement channels. Chile's public medicine procurement agency aggregated the requirements of public hospitals and primary-care providers and procured medicines for the public health network (Central de Abastecimiento del Sistema Nacional de Servicios de Salud, 2008). Private pharmacy chains instead purchased through commercial relationships with pharmaceutical laboratories and distributors. A 2001 competition ruling required manufacturers to publish medicine price lists and the criteria for volume and early-payment discounts. Because all three large chains qualified for the maximum volume discounts, the FNE characterized their acquisition (wholesale) costs as largely uniform and transparent (FNE, 2008). The same manufacturer lists included a common suggested retail price for each medicine, known as the PVPS (*precio de venta a público sugerido*). Each chain centrally set its own national retail price schedule. According to the FNE, the PVPS generally gave the chains a gross margin of 20–25% and provided a common

[5] The shares are the FNE's, citing IMS Health Chile. See Online Appendix A, "(b) Three chains setting one national price schedule."

[6] Online Appendix A, "(c) The FNE's account of prescription constraints."

benchmark for profitable retail pricing.[7]

2.2 From price war to coordinated increases

The PVPS was only a benchmark. Each chain could price below it. The price record passes through four episodes. First, through September 2006, the chains offered regular Monday and Thursday discounts that normally reversed within three days. These were temporary promotions rather than a sustained reduction in the price level.

Second, the temporary promotions became persistent cuts. Salcobrand stopped reversing its discounts on 22 October 2006, Cruz Verde did so on 23 November, and FASA lowered its prices more gradually. From then through mid-2007, the chains repeatedly undercut one another and average prices declined.[8] The chains used high-turnover branded medicines as loss leaders, trading losses on the discounted medicines for profits on other purchases during the same pharmacy visit (TDLC, 2012, consids. 42–43).

Third, Cruz Verde formalized the price war in August 2007 with a campaign that compared its prices on 685 high-turnover branded medicines with FASA's prices.[9] By publicizing the price difference, comparative advertising allowed the lowest-priced chain to attract shoppers from its rivals. FASA and Salcobrand responded by cutting prices on the advertised medicines to defend their market shares, pushing many prices below wholesale cost. FASA asked the Council for Advertising Self-Regulation and Ethics (*Consejo de Autorregulación y Ética Publicitaria*, CONAR) to stop the campaign. CONAR ruled against Cruz Verde on 7 September and upheld that ruling on 5 October. Both decisions were nonbinding.

Fourth, FASA obtained a binding advertisement ban from Santiago's 17th Civil Court on 6 November.[10] The binding ban ended the comparison campaign, removing the channel through which a chain could publicize its low prices and attract shoppers. FASA's chief

[7] Online Appendix A, "(d) The PVPS was a common, manufacturer-supplied benchmark," reproduces the FNE's description of the PVPS and the associated margin. The tribunal left this margin characterization unresolved.

[8] Online Appendix A, "(e) Temporary discounts became persistent cuts at the end of 2006," documents the break in the daily price record and reports the validation figure.

[9] Online Appendix A, "(f) Comparative advertising intensified the price war in August 2007," gives the complaint's description, the press record, and the uncontested finding that the price war compressed margins and produced below-cost sales.

[10] Online Appendix A, "(g) The advertisement ban became binding on 6 November 2007," documents the date and FASA's account of why further cuts stopped making sense (17º Juzgado Civil de Santiago, 2007). Figure 2 plots the relative-quantity response around an undercut.

executive consequently described further price cuts as no longer making sense, and the sustained price war ended at the same time. Individual price cuts continued. The chains then used a separate laboratory-mediated procedure to sequence price increases, beginning in December 2007. The FNE's requests for information reached the chains at the end of March 2008, after which new coordinated increases slowed. A few more coordinated increases occurred in April, followed by a pause. Coordinated increases resumed in July, while prices that had already been raised largely stayed at their higher levels.

Figure 1 places the institutional dates alongside the weekly average-price series. Panel (a) reports the legal and case chronology. Panel (b) shows the stable promotional period, the decline during persistent undercutting, and the price recovery during the cartel wave (9 December 2007–31 March 2008).

Figure 1: From price war to coordinated price increases

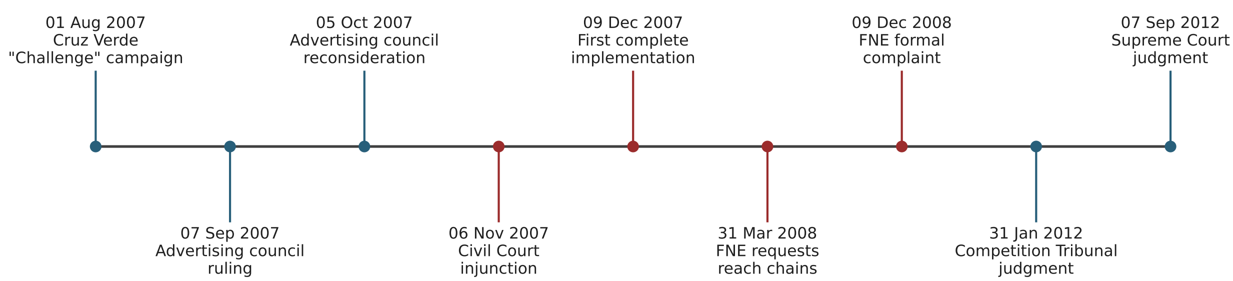

(a) Institutional timeline

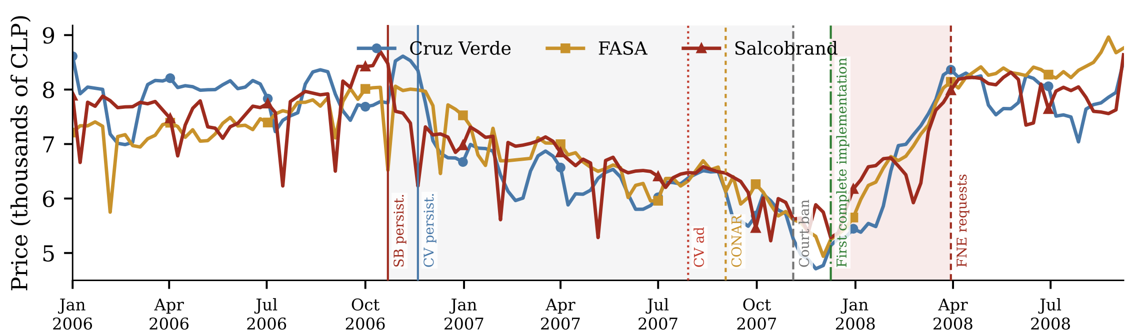

(b) Weekly average prices by chain

Note: Panel (b): quantity-weighted weekly prices across 222 medicines, in thousand CLP. Gray shading denotes the price war; red denotes the initial cartel wave (9 December 2007–31 March 2008). CV: Cruz Verde; SB: Salcobrand; "persist.": persistent cuts. Vertical lines mark the events discussed in the text.

2.3 Laboratory-mediated coordination

The evidence below traces how the chains established a common way to sequence and monitor price increases, then expanded it after the first successful batches. At first, they lacked this common procedure. Laboratory mediation resolved that coordination problem without removing the leader's uncertainty about whether rivals would participate in further increases. The price war left many medicines below wholesale cost. Restoring margins required one chain to raise first without knowing whether both rivals would follow. In its defence, Salcobrand stated that its new management sought to restore margins even at the risk of losing market share. Salcobrand offered to raise first, leaving FASA and Cruz Verde three or four days to detect and match each increase.[11]

Pharmaceutical laboratories became intermediaries in the one-two-three procedure. Their managers communicated proposed prices and dates, informed the other chains after the first chain moved, and checked whether all three had followed. The chains monitored compliance through price checks by their own or contracted staff and reports from the laboratories.[12] When either rival did not follow, the initiator typically returned to its previous price. A Bayer batch illustrates the sequence. Salcobrand raised the full list on Day One, and FASA and Cruz Verde each raised half on Day Two and the remainder on Day Three (TDLC, 2012, consid. 75).[13]

The case record contains conflicting accounts of who first proposed the procedure. FASA and Salcobrand attributed the proposal to the laboratories, whereas laboratory managers said that the chains asked them to contact rivals and coordinate dates. The TDLC read the record as showing that the push for coordination came from the pharmacy side. A Salcobrand email dated 19 December reported successful early batches for five products involving four laboratories and expected the procedure to be repeated with more products and laboratories.[14]

[11]Online Appendix A, "(h) Salcobrand's new management sought to restore margins," documents Salcobrand's stated margin strategy. "(k) The one-two-three procedure: Day One, Day Two, Day Three," reproduces the internal email describing its offer to move first.

[12]Online Appendix A, "(l) Monitoring: prior knowledge and quotation intensity," documents these monitoring channels; see also FNE, 2008, para. 96, p. 35.

[13]Online Appendix A, "(j) What the laboratories transmitted," reproduces the complaint's description, one laboratory email, and one laboratory-manager statement.

[14]Online Appendix A, "(i) Who proposed the coordination," presents the conflicting accounts and the tribunal's reasoning. The Salcobrand email of 19 December 2007 is reproduced in "(n) The agreement expanded after successful early batches"; see also TDLC, 2012, consids. 95 and 97. Online Appendix Section B.4, "Laboratory batches," documents the

The three chains' daily transaction-price paths record the first complete three-chain implementation on 9 December 2007. I use this observed implementation to date the model's transition to a commonly understood procedure. Under the weekly information transition in Section 5.2, informed play begins the following week. The FNE's requests for information reached the chains at the end of March 2008. The formal notice to Salcobrand was Oficio No. 419. The FNE filed its formal complaint on 9 December 2008, and the TDLC later sanctioned Cruz Verde and Salcobrand for coordinating prices on at least 206 medicines (TDLC, 2012, consids. 191, 205–206).[15]

The chains argued that the price increases reflected independent responses to a price leader rather than collusion. The tribunal acknowledged that the price patterns alone could support either explanation. It found collusion after considering the emails and testimony alongside the price evidence.[16]

2.4 Data and event construction

(i) Prices and quantities. I use a processed daily panel covering 222 medicines at Cruz Verde, FASA, and Salcobrand from 2006 through 2008. Each chain–drug–day observation contains a revenue-weighted average transaction price and units sold. Separate listed-price series cover the same panel. I use transaction prices to date events and to measure realized prices, markups, price tiers, and demand. Salcobrand's wholesale costs, available from November 2007 onward, come from evidence submitted to the TDLC.[17]

(ii) Event panel. I classify persistent drug-level price plateaus into three tiers. Tier 0 is the below-cost tier, with quantity-and-cost-weighted cost-relative markups of 9–12% below wholesale cost across chains. Tier 1 is the margin-restoration tier, the first common price that restores a positive retail margin, with markups of 22–25% above wholesale cost. Tier 2 is the rent-extraction tier, the next common price above Tier 1, with markups of 43–48%

timing using the three chains' daily transaction-price paths.

[15]Online Appendix A, "(o) Scope, chronology, and outcome," and "(p) Communication after the investigation began." The judgment dates Oficio No. 419, addressed to Salcobrand, to 3 April 2008; testimony places the FNE's requests at the end of March. The judgment's "at least 206 medicines" is a legal minimum for December 2007–March 2008, not the event count defined below.

[16]Online Appendix A, "(q) The alternative explanations on the record," presents the competing accounts and the tribunal's assessment; see also TDLC, 2012, consids. 159, 165–166, 192.

[17]Online Appendix Section B.1, "Data construction," describes the data and sample selection.

above wholesale cost. These markups use wholesale cost as the denominator, $(p-c)/c$, rather than the retail-price margin defined above. Table 2 reports the chain-specific values. For each medicine I compute one benchmark price per tier, common across the three chains, by averaging their weekly transaction prices over the tier's window. Tier 0 uses the price-war period. Tiers 1 and 2 use stable periods beginning two weeks after the corresponding first successful increase. Where a benchmark price is missing, I use the median Tier 1-to-Tier 0 price ratio across medicines for Tier 1. For Tier 2, I use a markup regression estimated on medicines with an observed stable Tier 2 window for Tier 2.[18] I construct the event panel from the three chains' daily transaction-price paths. The detection procedure flags sustained changes relative to a chain's recent price and its rivals' prices. Coordinated-increase candidates are then checked for their date, initiator, rival response, and persistence. The panel records successful and unsuccessful increases and price cuts before and after coordination begins.

An increase event is one medicine-level attempt, including the initiating chain's price change and the rivals' responses. It is successful when all three chains reach their chain-specific prices within the same higher tier. The attempt is unsuccessful if the third chain does not follow. For each medicine, I record the first time its price crosses each adjacent tier boundary. A direct Tier $0 \rightarrow 2$ increase is therefore one event, but the observed boundary counts record it as the medicine's first crossing of both the Tier $0 \rightarrow 1$ and Tier $1 \rightarrow 2$ boundaries.

(iii) Documentary record. I use the FNE complaint, the TDLC decision, the CONAR ruling, and the civil-court injunction to reconstruct the institutional timeline and communication among the chains.[19]

3 Descriptive Evidence

This section documents four patterns used in the demand system and dynamic game. Temporary discounts became persistent, the quantity response to an undercut changed after the advertisement ban, and increase attempts became organized by laboratory and date. The movement from the below-cost price to the first common price with a positive retail margin (Tier $0 \rightarrow 1$) was faster than the subsequent movement from that positive-margin price to the next higher common price (Tier $1 \rightarrow 2$). The model takes the start of the price war as

[18] Online Appendix Section B.2, "Construction of price tiers," explains how I calculate the benchmark tier prices.

[19] Online Appendix A, "(a) Sources and evidentiary conventions," identifies the document supporting each institutional fact.

given.

3.1 From temporary discounts to persistent low prices

For most of 2006, discounts at Cruz Verde and Salcobrand reversed within days. Among January–September markdowns of at least 15%, 96.4% at Salcobrand and 95.2% at Cruz Verde returned to at least 95% of the pre-markdown price within three days. This pattern changed late in the year. Salcobrand cut 129 medicines by a median of 21% on 22 October, and only one returned to that threshold the next day. Cruz Verde cut 164 medicines by a median of 24% on 23 November, and none returned the next day. By late November, the share of medicines with weekly median prices at least 10% below their January–September median of weekly median prices exceeded one-half at both chains in each of four consecutive weeks (Online Appendix Figure 1, Panel (c)).[20]

3.2 The quantity response to an undercut

I define an undercut event as a transaction-price cut of at least 15% from a stable seven-day plateau that makes the initiator newly cheapest. I compare the initiator's $\log(1+q)$ with the mean $\log(1+q)$ of its two rivals around the event. Here, q is daily quantity sold. I subtract the mean of this difference over the seven days before the cut. Figure 2 plots this relative log quantity from seven days before through seven days after the cut. Before the advertisement ban, the response on the cut date is large but reverses on days one and two. After the ban, the cut-date response is small and the increase appears on days one and two. These event paths motivate the demand specification in Section 4, which allows the gain from belonging to the low-price group to differ across the two regimes for the same medicine.

[20] The markdown counts require consecutive-day cuts of at least 15%. Prices are observed on both days for 189 Salcobrand medicines and 209 Cruz Verde medicines. Contemporary SERNAC evidence also documents regular weekday and card-based promotions at both chains: Servicio Nacional del Consumidor, "SERNAC denuncia publicidad engañosa de farmacias," 6 June 2006, https://www.sernac.cl/portal/619/w3-article-864.html.

Figure 2: Relative quantity around an undercut

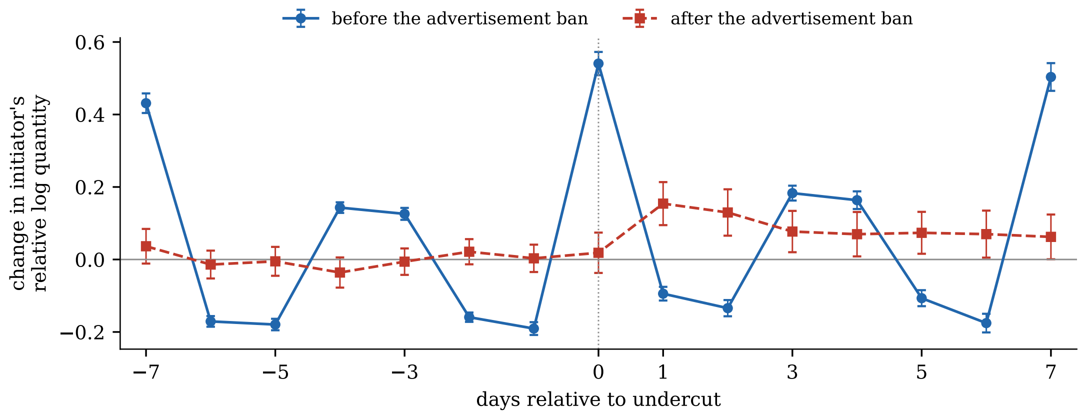

Note: Pre-ban: 4,947 events through 31 August 2007; post-ban: 373 events from 6 November 2007. Relative log quantity is normalized to its days −7 to −1 mean. Bars show 95% confidence intervals with standard errors clustered by medicine. The demand-estimation post-ban sample starts on 12 November.

3.3 Organizing coordinated price increases

The three chains' daily transaction-price paths show that observed increase attempts became more concentrated by laboratory and date before the first successful three-chain implementation. I group attempts out of Tier 0 when they concern medicines supplied by the same laboratory on the same date. From 2 September through 5 November 2007, the 25 unsuccessful attempts form 24 such groups across 22 dates. From 6 November through 8 December, 42 unsuccessful attempts form 28 laboratory–date groups.[21] Table 2 reports the composition and outcomes from 6 November onward. This increase in grouping by laboratory and date disciplines the laboratory review process in Section 5. Figure 3 places these attempts in the full episode. An unsuccessful attempt is an increase not matched by both rivals. The initiator may revert or remain at the higher price.

[21]Online Appendix Section B.4, "Laboratory batches," reports both windows.

Figure 3: Coordination events over the episode

(a) Events by tier transition

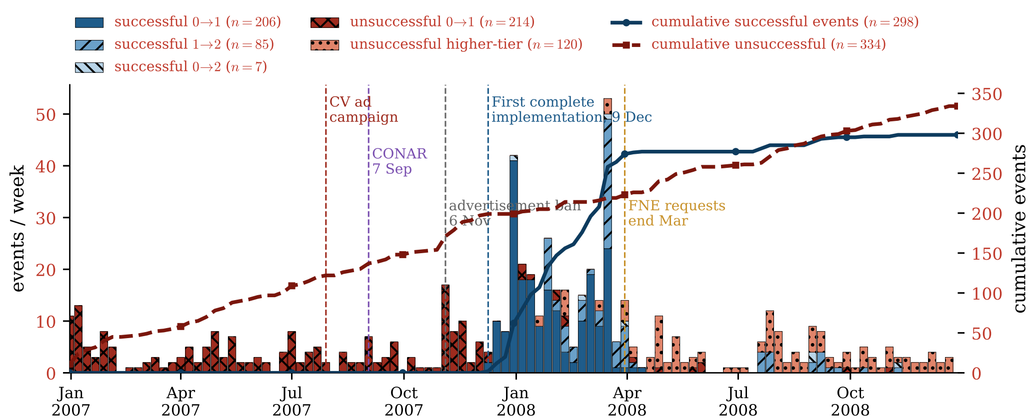

(b) Initiations by chain and period

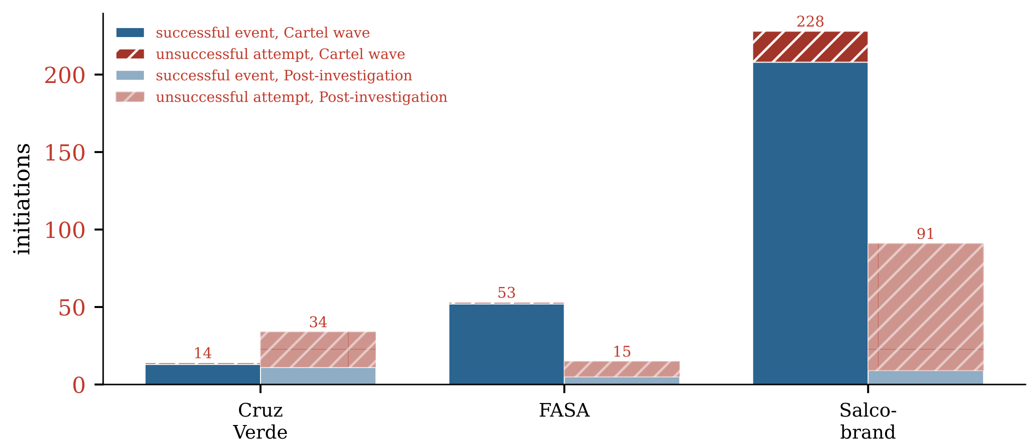

Note: Panel (a): daily-price events, 31 December 2006–3 January 2009. Panel (b): cartel wave (9 December 2007–31 March 2008) versus 1 April–31 December 2008. Each event is one medicine-level initiation attempt; Panel (b) assigns it to the initiating chain. Bar labels count successful and unsuccessful initiations combined. "Unsuccessful higher-tier" combines Tier 1 and Tier 2 attempts. CV denotes Cruz Verde.

Panel (a) of Figure 3 shows repeated unsuccessful attempts before the first complete implementation on 9 December 2007. Successful Tier 0 $\rightarrow$ 1 events then spread through March 2008, while Tier 1 $\rightarrow$ 2 events began on 27 January 2008. A few further increases succeeded in April after the FNE's requests reached the chains. After the last of these April increases, successful increases stopped until late July, although unsuccessful attempts continued.[22] These later events describe the continuation of coordination beyond the dynamic estimation period.

Panel (b) shows that Salcobrand initiated 208 of the 273 successful events during the cartel wave. Initiation was less concentrated from 1 April through 31 December 2008. Salcobrand initiated 9 successful events, Cruz Verde 11, and FASA 5.[23]

After the FNE investigation began, Cruz Verde and FASA shifted from emails and telephone calls to face-to-face contact. Salcobrand retained its ordinary communication channels.[24] The later price pattern retained some laboratory grouping. Three of the 19 Tier 1 $\rightarrow$ 2 events, excluding direct Tier 0 $\rightarrow$ 2 events, after 31 March 2008 occurred on 25 July 2008 for Pfizer. This formed one same-day, same-laboratory cluster.[25]

To characterize which medicines entered coordination earlier, I regress the calendar date of each drug's first successful event out of Tier 0 on its pre-ban price coefficient, a drug-specific descriptive demand-slope measure constructed from pre-ban price and quantity variation, mean pre-ban revenue, mean positive gap between wholesale cost and price, and drug type (chronic or acute crossed with prescription or over-the-counter status). These characteristics explain 10.6% of the variation in these event dates. Adding laboratory fixed effects raises the R^2 to 63.6% (Table 1), so the observed ordering is strongly associated with laboratory membership. Predetermined product attributes explain little of the remaining order within laboratories.[26]

Across the 36 laboratories, the correlation between mean log pre-ban product revenue and mean date of the first successful event out of Tier 0 is -0.22. Laboratories with higher-revenue medicines entered earlier on average.[27]

---

[22] Online Appendix Section B.4, "After the investigation," separates these attempts from cuts away from an established coordinated price.

[23] Online Appendix Section B.4, "Leadership by chain and period," reports the full decomposition.

[24] Online Appendix A, "(p) Communication after the investigation began," reports that Cruz Verde and FASA moved from emails and telephone calls to face-to-face contact, while Salcobrand continued to use its ordinary channels (TDLC, 2012, consid. 88).

[25] Online Appendix Section B.4, "After the investigation," reports the construction.

[26] See Online Appendix Section B.4, "Within-laboratory coordination order," Table 3.

[27] The correlation aggregates 210 coordinated medicines to the laboratory level. Drug-level

Table 1: Laboratory membership and timing of the first successful event out of Tier 0

Specification ($N = 210$)	R^2
Price coefficient, pre-ban revenue, below-cost gap, and drug-type fixed effects	0.106
+ laboratory fixed effects	**0.636**

Note: The outcome is the calendar day of each medicine's first successful event out of Tier 0. The sample contains the 210 medicines whose first successful event out of Tier 0 occurred by 30 June 2008. The pre-ban price coefficient is a drug-specific descriptive demand-slope measure; Section 4 estimates the demand system jointly across medicines. Pre-ban revenue and the below-cost gap are measured in thousands of CLP.

3.4 The movement beyond restored margins

Table 2 reports quantity-and-cost-weighted markups of 9–12% below cost in Tier 0, 22–25% above cost in Tier 1, and 43–48% above cost in Tier 2. Figure 4 compares the cross-medicine mean of Salcobrand's wholesale costs for the 221 medicines with observed costs against the quantity-weighted transaction price across three chains and 222 medicines. From the common 1 November 2007 base to 1 May 2008, the retail-price index rose 43% while the wholesale-cost index rose 2.2%.[28]

Figure 4: Wholesale cost and retail price

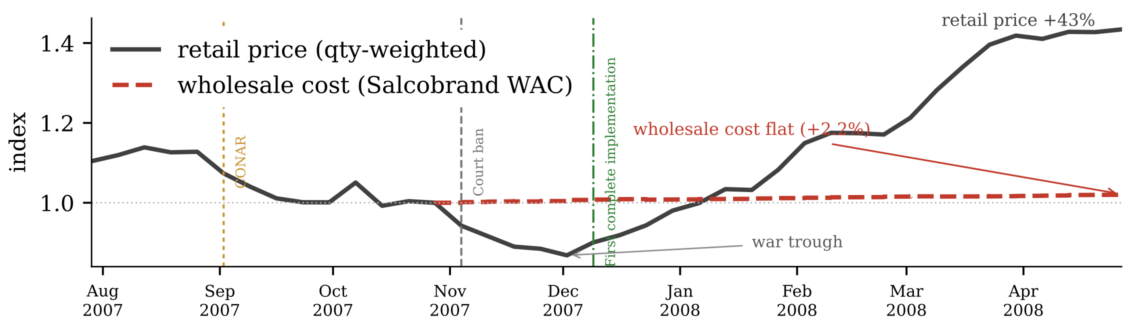

Note: Both indices equal one on 1 November 2007. Dashed: wholesale cost; solid: retail price. WAC denotes wholesale acquisition cost; "Court ban" denotes the advertisement ban. Samples and weighting are described in the text.

pre-ban revenue is mean daily chain–drug revenue before the advertisement ban; timing is the laboratory mean of each medicine's first successful event date out of Tier 0.

[28]Online Appendix Section B.3, "Monthly losses on below-cost medicine sales," reports the associated descriptive peso shortfall and its cost-data limitation.

Table 2: Cost-relative markups by tier and chain

	Tier 0 Below cost	Tier 1 Margin restoration	Tier 2 Rent extraction
Cruz Verde	−0.101	0.236	0.454
FASA	−0.117	0.220	0.435
Salcobrand	−0.089	0.249	0.476

Note: Each cell is a quantity-times-cost-weighted average of the cost-relative markup: $\sum_j Q_{ij}(p_j^z - c_j)/\sum_j Q_{ij}c_j$. Here p_j^z is medicine j's fixed benchmark transaction price in Tier z, for $z \in \{0, 1, 2\}$, c_j is its wholesale cost, and Q_{ij} is chain i's total quantity sold across the full demand panel, using the same quantity weights for all three tiers. The fixed benchmark tier prices and costs are common across chains within a medicine, so chain differences in this table come from quantity weights. The sample contains 222 medicines, each contributing once to each chain–tier cell. Online Appendix Section B.2, "Construction of price tiers", gives the tier construction and imputations.

Coordination continued after the Tier 0 → 1 movement had restored positive margins. From 31 December 2006 through 3 January 2009, the three chains' daily transaction-price paths record 206 successful Tier 0 → 1 events, 85 successful Tier 1 → 2 events, and seven direct Tier 0 → 2 events. These categories are mutually exclusive. A direct Tier 0 → 2 event crosses both adjacent tier boundaries. The first Tier 1 → 2 crossing occurred on 27 January 2008, 49 days after the first Tier 0 → 1 crossings on 9 December 2007. Over the full observation window, Salcobrand initiated 59 successful Tier 1 → 2 events.[29]

Margin restoration proceeded faster even after accounting for how many medicines were eligible for each movement and how long they remained eligible. I compare the 65 days beginning with the first success for each movement. The Tier 0 → 1 window is 9 December 2007–11 February 2008. The Tier 1 → 2 window is 27 January–31 March 2008. For Tier 0 → 1, medicines remain eligible until their first successful increase out of Tier 0. For Tier 1 → 2, they enter the eligible group the day after completing Tier 0 → 1 and leave after completing Tier 1 → 2. I divide the number of first transitions by the total time these medicines remain eligible. The resulting rates are 9.83 and 5.39 transitions per 100 medicine-weeks, respectively. Margin restoration occurred at 1.83 times the rate of rent extraction. Direct Tier 0 → 2 events are counted separately because they skip the waiting time in Tier 1.[30]

[29] Online Appendix Section B.4, "Leadership by chain and period", Table 4 reports successful events by initiator and the rates of movement between adjacent tiers.

[30] Online Appendix Section B.4, "Leadership by chain and period," Table 4, Panel B, reports the calculation.

Section 5 compares the observed diffusion paths with separately estimated models in which the leader's subjective weight either updates after rent-extraction attempts or remains fixed at one.

4 Demand

I estimate a demand system to measure how price differences and recent price changes affect purchases at each chain. These estimates translate the price changes described in Section 3 into changes in quantity and profit. Section 5 then uses these profits to estimate the dynamic pricing game.

4.1 Dynamic instrumental-variable (IV) demand

I estimate a simple-logit demand system that allows purchases to respond to current prices, low-price-group membership, and recent price changes. The low-price-group term captures the additional demand a chain attracts by pricing sufficiently below its rivals. The price-history terms allow a recent increase or cut to affect demand beyond its effect on the current price. They also capture an additional response when a chain raises its price and remains alone at the high price. These features connect demand to the price changes documented in Section 3.

I use a daily panel of recorded transaction prices and quantities. Let $j \in \{1, \ldots, J\}$ index medicines, $i \in \mathcal{I}$ index the three pharmacy chains, and t index days. For each medicine, consumers choose among the three chains and the outside good of not purchasing from any of them. I retain a medicine–day only when all three chains report positive prices and quantities.

Let q_{ijt} denote quantity and N_j the market size for medicine j. I set N_j equal to the maximum total daily quantity sold by the three chains on a day with complete observations, divided by 0.92, plus one unit.[31] The inside and outside shares are therefore

$$s_{ijt} = \frac{q_{ijt}}{N_j}, \qquad s_{0jt} = 1 - \sum_{i \in \mathcal{I}} s_{ijt}.$$

Consumer n's utility from purchasing medicine j at chain i on day t is $u_{nijt} = \delta_{ijt} + \varepsilon_{nijt}$, where δ_{ijt} is mean utility and ε_{nijt} is an idiosyncratic taste shock. I normalize the outside

[31] The maximum accommodates daily sales fluctuations. The divisor uses the three chains' combined market share reported in the Introduction, and the additional unit ensures a positive outside share.

good's mean utility to zero, so $u_{n0jt} = \varepsilon_{n0jt}$. With i.i.d. Type I extreme-value taste shocks, the logit shares imply $\log(s_{ijt}/s_{0jt}) = \delta_{ijt}$.

Mean utility depends on the transaction price p_{ijt}, low-price-group membership L_{ijt}, recent increases A^{+}_{ijt}, recent cuts A^{-}_{ijt}, and the unmatched high-price position H_{ijt}. I allow the price coefficient and the low-price-group coefficient to differ before and after the advertisement ban, while keeping the three price-history coefficients common across periods. Let Post_t equal one in the post-ban estimation window and zero in the pre-ban window. The estimating equation is

$$\begin{aligned} \log(s_{ijt}/s_{0jt}) = {} & \phi_{ij} + \zeta_t - \alpha_{\text{pre}} p_{ijt}(1 - \text{Post}_t) - \alpha_{\text{post}} p_{ijt} \text{Post}_t \\ & + \lambda_{\text{pre}} L_{ijt}(1 - \text{Post}_t) + \lambda_{\text{post}} L_{ijt} \text{Post}_t \\ & + \gamma_{\text{inc}} A^{+}_{ijt} + \gamma_{\text{cut}} A^{-}_{ijt} + \gamma_H H_{ijt} + \xi_{ijt}. \end{aligned} \tag{1}$$

Here ϕ_{ij} and ζ_t are chain–medicine and common-day fixed effects, and ξ_{ijt} is the remaining demand shock. Price is measured in thousands of Chilean pesos. The positive price coefficients α_{pre} and α_{post} enter with a minus sign. The λ coefficients measure the additional mean utility associated with low-price-group membership.

A chain belongs to the low-price group when its price is at least five percent below the highest chain price. Writing $p^{\max}_{jt} = \max_{k \in \mathcal{I}} p_{kjt}$, I define

$$L_{ijt} = \mathbf{1}\left\{ p_{ijt} < p^{\max}_{jt} \text{ and } \frac{p^{\max}_{jt} - p_{ijt}}{p^{\max}_{jt}} \geq 0.05 \right\}.$$

Up to two chains can therefore belong to the low-price group for the same medicine on the same day.

The recent-increase and recent-cut variables accumulate price changes over the current day and the preceding six calendar days. For consecutive daily prices, let $\Delta p_{ijt} = \log(p_{ijt}/p_{ij,t-1})$, and set it to zero when consecutive daily prices are unavailable. With ℓ denoting the number of days before t,

$$A^{+}_{ijt} = \sum_{\ell=0}^{6} \max\{\Delta p_{ij,t-\ell}, 0\}, \qquad A^{-}_{ijt} = \sum_{\ell=0}^{6} \max\{-\Delta p_{ij,t-\ell}, 0\}.$$

Finally, H_{ijt} captures a recent increase that leaves the chain alone at a sufficiently high price. Let $\text{gap}_{ijt} = \max\{0, \log p_{ijt} - \log \max_{k \neq i} p_{kjt}\}$ when $\max_{k \neq i} p_{kjt} \leq 0.95 p_{ijt}$, and zero otherwise. The threshold measures the highest rival's discount relative to the chain's own price. I define $H_{ijt} = \text{gap}_{ijt} \mathbf{1}\{A^{+}_{ijt} > 0\}$. This term becomes zero when a rival matches the price, even if the recent increase still contributes to A^{+}_{ijt}. Unlike L_{ijt}, the three price-history variables measure magnitudes rather than binary events. I use these same definitions in both demand regimes and recompute them under counterfactual price paths using only prices available through day t.

4.2 Identification and estimation

I use national listed prices as instruments for transaction prices. A listed price is the chain's centrally chosen national pre-discount schedule, whereas transaction prices incorporate the discounts and promotions applied to customers (TDLC, 2012, consids. 27–31). The case record also describes list prices and manufacturers' suggested prices as strategic benchmarks (FNE, 2008, para. 96). Transaction prices can respond to transaction-stage demand shocks through discounts and promotions, making them endogenous. My maintained identifying assumption is that changes in the national listed-price schedule, including the schedule against which weekly discounts are set, are predetermined strategic pricing choices and do not respond to the residual consumer-demand shock ξ_{ijt}. Conditional on chain–medicine and day fixed effects and the other demand controls, listed prices affect purchases through transaction prices. The institutional timing of the national schedule therefore motivates the exclusion restriction. I interact both prices with the pre- and post-ban indicators to estimate the two price coefficients.

I assume that the two low-price-group interactions and the recent price-history variables A^{+}_{ijt}, A^{-}_{ijt}, and H_{ijt} are uncorrelated with the remaining demand shock after accounting for chain–medicine and day fixed effects. These fixed effects absorb persistent differences in demand across chain–medicine pairs and common daily shocks. I estimate all seven demand coefficients jointly by two-stage least squares and cluster standard errors by medicine.[32]

The estimation sample covers 1 January 2006–2 September 2007 before the advertisement ban and 12 November 2007–31 December 2008 afterward. I exclude the intervening period, which covers the advertising rulings and extends through the first calendar week of the binding ban. Listed prices are available for 589,624 chain–medicine–day observations, or 95.5 percent of the eligible demand sample.

Table 3 reports the demand estimates. The price coefficient changes little after the advertisement ban, falling from 0.034 to 0.028. The difference is −0.006, with a standard error of 0.003, while the mean own-price elasticity changes from −0.353 to −0.314. These estimates indicate a modest change in consumers' response to transaction prices.

The gain from belonging to the low-price group falls much more sharply. Its coefficient declines from 0.207 before the ban to 0.062 afterward, a reduction of 70 percent. The differ-

[32]The first-stage partial R^2 is 0.953 before the ban and 0.965 afterward; the corresponding medicine-clustered F-like statistics for the excluded instruments are 7,962 and 9,448.

Table 3: IV demand estimates by regime

	Pre-ban	Post-ban
Panel A. Regime-specific demand responses		
Price coefficient $\hat{\alpha}$	0.034***	0.028***
	(0.006)	(0.005)
Low-price-group gain $\hat{\lambda}$	0.207***	0.062***
	(0.016)	(0.014)
Mean own-price elasticity	−0.353	−0.314
Panel B. Common price-path responses		
Seven-day cumulative price-increase magnitude $\hat{\gamma}_{\text{inc}}$	−0.138**	
	(0.057)	
Seven-day cumulative price-decrease magnitude $\hat{\gamma}_{\text{cut}}$	−0.040	
	(0.047)	
Unmatched high-price-gap magnitude $\hat{\gamma}_H$	−1.013***	
	(0.101)	

Notes. Standard errors clustered by medicine are in parentheses. The Panel B coefficients multiply unscaled log price-change and log price-gap magnitudes. *, **, and *** denote significance at the 10, 5, and 1 percent levels. Mean own-price elasticity is the unweighted mean of $-\hat{\alpha}_d p_{ijt}(1 - s_{ijt})$ evaluated at observed prices and shares in regime d.

ence is −0.145, with a standard error of 0.018. Holding current prices, recent price histories, and fixed effects constant, low-price-group membership raises purchase odds relative to the outside good by 23.0 percent before the ban and 6.4 percent afterward. The main change in estimated demand is the smaller additional gain from offering a sufficiently low price.

Recent price changes also affect demand after controlling for current prices and low-price-group membership. The recent-increase coefficient is −0.138 (standard error 0.057). The recent-cut coefficient is −0.040 (standard error 0.047). It provides no clear evidence of an additional response to recent cuts. A recent increase carries a further demand loss when the chain remains alone at the high price. The coefficient on H is −1.013 (standard error 0.101). This additional loss grows with the log price gap between the chain and its highest-priced rival.

4.3 Demand-implied quantity and profit effects

A price change affects profit through both sales of the medicine and complementary basket profit associated with the resulting change in medicine demand. I use the demand estimates to calculate these two effects for an unmatched price increase and a unilateral five-percent cut. I treat each unit of medicine demand as one customer-equivalent unit and let μ denote complementary basket profit per customer-equivalent unit. Figure 5 shows how the profitability of each action varies with μ under pre- and post-ban demand.

I use price changes available in Section 5. The first increase starts from the observed initial price vector and raises one chain to Tier 1; I include only chains whose initial price is below Tier 1. The second starts at Tier 1 and raises one chain to Tier 2. Rivals hold their prices for seven days, and the unmatched raiser returns to its starting price next week. The cut is the first ordinary cut: one chain lowers its own initial price by five percent for seven days. Next week, all three chains charge 95 percent of their respective initial prices, the model's initial price-war state. Initial prices need not be equal across chains.[33]

The figure evaluates each increase or cut along the fixed two-week price path described above. This fixed path isolates how the estimated demand system changes the payoff from each action.

The calculation uses Section 5.1's one-off payoff. It includes seven days at the action prices plus the discounted seven-day price-history correction caused by the next week's price change. It excludes continuation values. First-week changes are measured against holding the starting prices. The next-week correction compares demand at the next price vector with and without price history; it does not include the full lasting effect of lower prices. Quantity percentages divide the first-week change plus this correction by baseline sales over the two weeks, evaluated at each week's price vector without price-history effects.

Let a denote the increase or cut. Write $\Delta q^0_{ij}(a)$ and $\Delta\pi^{0,\text{medicine}}_{ij}(a)$ for the first-week changes in daily quantity and medicine profit, and $\Delta q^1_{ij}(a)$ and $\Delta\pi^{1,\text{medicine}}_{ij}(a)$ for the second-week effects just described. Each medicine-profit term is the change in $q(p - c_j)$, where c_j is medicine j's wholesale cost. Adding profit from other purchases gives

$$\Delta\Pi^{\text{event}}_{ij}(a;\mu) = 7\left[\Delta\pi^{0,\text{medicine}}_{ij}(a) + \mu\Delta q^0_{ij}(a)\right] + \beta 7\left[\Delta\pi^{1,\text{medicine}}_{ij}(a) + \mu\Delta q^1_{ij}(a)\right], \quad (2)$$

where β is the weekly discount factor used in Section 5. Event profit is measured in thousand CLP, and μ in thousand CLP per customer-equivalent unit. The figure varies μ. Its value is estimated in the dynamic game.

[33] Online Appendix Section B.2, "Construction of price tiers," reports the price construction.

Figure 5: Demand-implied quantity and profit effects of unilateral price changes

(a) Unmatched-increase quantity loss

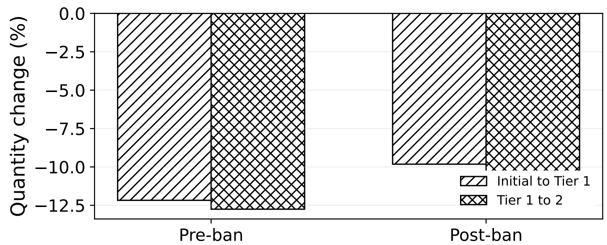

(b) Initial-cut quantity gain

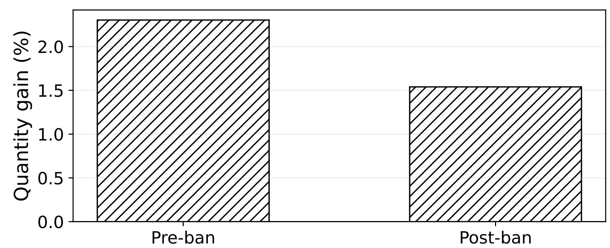

(c) Unmatched-increase profit gain/loss

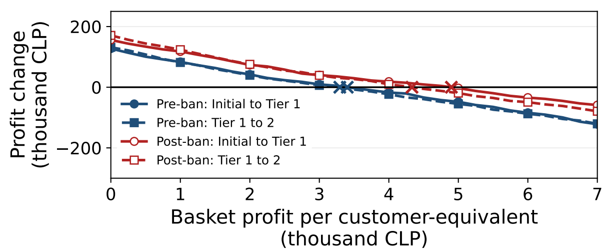

(d) Initial-cut profit gain/loss

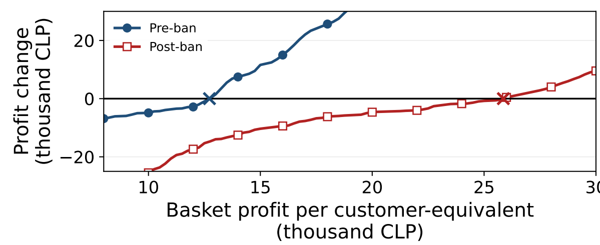

Notes. Medians across medicine–chain pairs: 467 for initial-to-Tier 1 increases; 666 for Tier 1 → 2 increases and cuts. Quantity changes are percentages of the two-week baseline; event profits are in thousand CLP. Complementary basket profit μ is in thousand CLP per customer-equivalent unit. Calculations follow Equation (2).

Panel A reports a median quantity loss of 12.2 percent before the ban and 9.8 percent afterward for an eligible initial-to-Tier 1 increase. The Tier 1 $\rightarrow$ 2 losses are 12.8 and 11.2 percent. Panel B reports a median gain of 2.3 percent before the ban and 1.5 percent afterward for the same initial five-percent cut. All these percentages use the two-week event denominator.[34]

Panel C shows how lost basket profit changes the return to an unmatched increase. The initial-to-Tier 1 profit curve crosses zero at $\mu = 3.40$ thousand CLP before the ban and 4.90 afterward; the Tier 1 $\rightarrow$ 2 crossings are 3.29 and 4.33. Panel D shows the opposite tradeoff for a cut. Additional basket profit can cover the lost medicine margin. The initial-cut curve crosses zero at $\mu = 12.73$ thousand CLP before the ban and 25.86 afterward. These are one-off payoff comparisons, not predictions of the optimal action. In Section 5, cutting after verification also starts the price-war continuation. That future loss can deter a cut even when its one-off payoff is positive. [35]

[34] For a separate seven-day comparison of predicted and observed losses following unmatched price increases, see Online Appendix Section C.3, "Observed outcomes around price changes."

[35] Online Appendix Section C.1, "Daily quantity and profit," reports profit thresholds

5 Dynamic Pricing Game

This section develops a dynamic pricing game to model the transition from price war to coordinated price increase. It organizes the transition around three mechanisms. First, the advertisement-ban mechanism operates through the estimated demand system: the post-ban regime has a smaller fitted reward from being in the low-price group and thereby changes the payoff from undercutting. Second, the procedure-verification mechanism allows initially uninformed followers to learn a repeatable one-two-three price-increase procedure, which supports coordinated margin restoration across medicines (Tier $0 \to 1$). Third, after verification, the model allows a candidate leader's subjective weight on sequential follower choices, relative to raising alone, to respond to earlier rent-extraction outcomes (Tier $1 \to 2$). The **Adaptive Confidence model** updates this weight after each complete or incomplete selected rent-extraction proposal using a Beta-shaped outcome-updating rule. In this model, actual followers make their own sequential choices. The separately estimated **Unit Confidence model** instead fixes the leader's subjective weight at one after verification, retaining the same sequential follower choices, and provides the benchmark.

The demand estimates determine quantities, which enter flow payoffs through medicine margins and complementary basket profit. The same μ measures complementary basket profit in every block. What changes after the advertisement ban is the demand generated by a low price. Observed laboratory dates and counts determine scheduled reviews in the laboratory-mediated coordination game (Game 2). Verification changes what firms know about the price-increase procedure, while followers continue to make their own participation decisions. Both specifications use the same demand estimates and review rules but estimate their economic parameters separately. Margin restoration can follow once the procedure is verified. Rent extraction goes further and requires the leader to assess whether rivals will participate.

5.1 Environment, states, timing, and flow payoffs

Model setup. Let $i \in \mathcal{I} = \{\mathrm{CV}, \mathrm{FASA}, \mathrm{SB}\}$ index pharmacy chains, $j \in \mathcal{J} = \{1, \ldots, J\}$ index medicines, and $t \in \mathcal{T}^{\mathrm{obs}} = \{52, \ldots, 117\}$ index Sunday–Saturday model weeks. Here t indexes weeks. In the demand section it indexes days. Thus, $\mathcal{I}$ contains Cruz Verde, FASA,

for the price-change week, excluding next-week price-history effects and later continuation values. Section C.2, "Demand specifications and quantity effects of model actions," compares event quantity losses from increases and gains from cuts across demand specifications. Online Appendix Table 8 shows that the post-ban cutoff is higher across alternative price-history definitions.

and Salcobrand, $J = 222$, and $\mathcal{T}^{\text{obs}}$ runs from 31 December 2006–6 January 2007 ($t = 52$) through 30 March–5 April 2008 ($t = 117$).

I divide this period into two games because the binding advertisement ban marks an exogenous regime change in the model. It changes both payoffs and how firms receive opportunities to raise prices. The decentralized price-adjustment game (Game 1) covers 31 December 2006–3 November 2007 ($t = 52, \ldots, 95$). During this game, each eligible medicine may receive a weekly price review. Within Game 1, the model switches to calibrated campaign cut depths and a separately estimated review frequency in 5–11 August 2007 ($t = 83$). The earlier block ends with 29 July–4 August 2007 ($t = 82$), and the campaign block ends with 28 October–3 November 2007 ($t = 95$). I use 5 August as the model cutoff.[36] The laboratory-mediated coordination game (Game 2) covers 4 November 2007–5 April 2008 ($t = 96, \ldots, 117$). The binding advertisement ban took effect on 6 November 2007, during the first week of Game 2. From Game 2 onward, the model uses post-ban payoffs and closes the independent Game 1 review process; laboratory dates and product counts determine scheduled reviews in Game 2. The model assigns both changes to the 4–10 November 2007 week ($t = 96$).

In the forward simulation, the observed laboratory dates and medicine counts determine the scheduled reviews. At each date, the model randomly selects up to the observed count from the eligible medicines without replacement. The stationary value calculation uses the corresponding average review opportunity. Random assignment and cut-blocked reviews can leave some medicines without an unblocked opportunity to restore margins within the simulation period. I therefore add supplementary Tier $0 \to 1$ reviews so that the allocation of review slots does not by itself exclude these medicines from a verified opportunity. These reviews provide an opportunity, not a guaranteed increase. There are no supplementary Tier $1 \to 2$ reviews.[37] These reviews determine when a medicine can be considered for an increase. Verification determines whether the followers know the procedure for coordinating that increase. A candidate leader is a chain facing the choice to raise first. The first chain to accept becomes the leader. Before verification, only Salcobrand faces this choice. After verification, any of the three chains can do so.[38]

Before verification, the leader and followers solve different subjective games under the

[36]The record dates the campaign only to August 2007. The model uses 5 August 2007 because it is the first complete model week beginning in August; it is not an observed campaign-start date.

[37]Online Appendix Section D.3, "Forward simulation," specifies the timing and eligibility of supplementary reviews.

[38]See Online Appendix Section B.4, "Laboratory batches."

same post-ban demand system. Salcobrand knows the laboratory procedure and solves Game 2. Each uninformed follower instead solves a decentralized game with the structure of Game 1's campaign block (Block 2), recomputed using post-ban demand. It expects the other chains, including Salcobrand, to follow that decentralized game and does not anticipate the laboratory procedure. In each unverified week, the followers become informed with probability η_0. Once verified, they switch permanently to Game 2 and, in both specifications, choose whether to follow according to their own incentives. In Adaptive Confidence, a candidate leader evaluates a rent-extraction proposal by weighting the value of sequential follower choices against the value of raising alone. Complete and incomplete attempts update this subjective weight. Under the Unit Confidence model, the candidate leader is assigned a subjective weight of one on the sequential follower outcome law after verification; each follower still chooses whether to join.

For drug j, let $z_{jt} \in \{I, 0^{(0)}, \ldots, 0^{(K)}, 1, 2\}$ denote its price state and $\mathbf{p}_j^{z_{jt}} = (p_{ij}^{z_{jt}})_{i\in\mathcal{I}}$ its persistent prices at the three chains. State I is the initial-price state, states $0^{(0)}, \ldots, 0^{(K)}$ are the price-war levels, and states 1 and 2 are the two coordinated-price tiers. K is the last index in the price-war grid, which starts at zero. At these prices, consumers choose among the three chains and the outside good. I define the transition rules after the within-week timing.

At the start of week t, let $G_t = (R_t, S_t^r, F_t^r)$ denote the public information common across drugs. The superscript r denotes rent extraction. The indicator R_t records whether the chains have verified the price-increase procedure. The counts S_t^r and F_t^r record completed and incomplete rent-extraction attempts made before week t. Equation (8) maps these counts, the initial subjective weight m_0^r, and prior strength ν^r (the effective number of outcomes represented by the prior) into the candidate leader's subjective weight m_t^r on the value of sequential follower choices relative to raising alone. This weight is not the probability that both rivals ultimately follow, since sequential choices can also produce incomplete participation. Thus, m_t^r is derived from G_t rather than included as a separate state variable. Because G_t is common across drugs, verification or the outcome of an attempt for one drug can affect later choices for another. For a drug outside punishment, I write the economic state as

$$x_{jt} = (z_{jt}, G_t). \tag{3}$$

Punishment is a separate continuation, with value $V_{ij,t}^{\text{pun}}(z_{jt})$, in which the drug has no future increase opportunities. Whenever a transition enters punishment, this value replaces the ordinary continuation value $V_{ij,t}(x_{jt})$; the same price state outside punishment retains its increase opportunities. Section 5.2 defines the punishment transitions. Calendar time t

determines which game and policy apply, but it is not part of the economic state.[39]

Weekly timing. For medicines outside punishment, price decisions occur at reviews, except that in Game 2 before verification, followers can also cut outside reviews. The simulation conditions these additional cuts on their observed weekly totals.[40] Punished medicines follow their cut-only continuation each week, without waiting for a review. Game 1 draws each eligible product with a common weekly probability. Game 2 uses the laboratory dates and their product counts. At a review, the three chains first choose cut or hold simultaneously. Any cut lowers the drug's price state in the following week and ends its decisions for the current week. If all three chains hold, the increase stage follows. Game 1 uses simultaneous increase-or-hold choices. Game 2 uses the candidate-leader and, when applicable, follower choices. Incomplete participation changes current prices only. complete participation advances the drug's persistent price tier. A cut after verification instead starts the punishment continuation and removes the drug from future increase opportunities. I place every price action on the first day of its model week. The model does not use the action's calendar day within the week.

Figure 6 summarizes the within-week sequence.

[39]I solve each block as a stationary infinite-horizon game. The value calculation approximates learning on a grid; the simulation keeps the learning counts and chain-specific prices. Online Appendix Sections D.2 and D.3 explain the solution and simulation, respectively.

[40]Online Appendix Section D.3, "Forward simulation," describes how these totals enter the simulation.

Figure 6: Within-week timing of the dynamic game

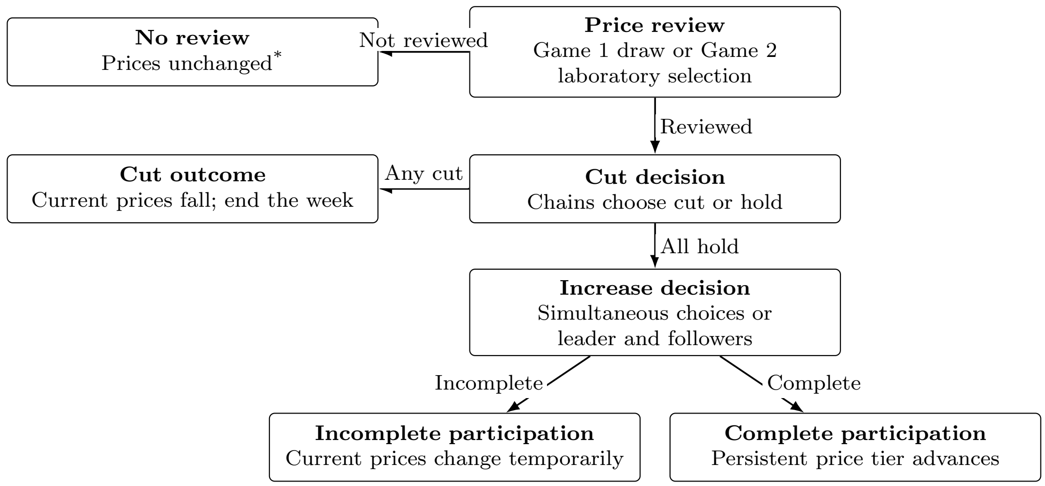

Note: Both games follow this sequence. *Before verification in Game 2, followers can also cut without a review; these cuts are conditioned on observed weekly totals. A cut after verification starts punishment and ends future increase opportunities.

Flow payoffs. The post-ban demand estimates reduce the customer-attraction benefit of undercutting in the flow payoffs below.

Let $\mathbf{p}_j^{\text{old}} \equiv \mathbf{p}_j^{z_{jt}}$ be the opening price vector. Let χ_{jt} denote the specified proportional cut depth for drug j in the current block. A cut action selects the corresponding lower price, so its depth enters the feasible action prices as well as the next-week state. Block 1 uses the ordinary five-percent cut. Block 2 allows drug-specific cut depths calibrated to observed price paths, and the post-ban games use the expanded cut-price space. The weekly choice maps it into one of three action-week price scenarios. If all chains hold at both the cut and increase stages, $a^c = a^u = \mathbf{0}$ and the price vector remains $\mathbf{p}_j^{\text{old}}$. If at least one chain cuts, the binary profile a^c identifies the cutters and the price vector is $\mathbf{p}_j^{a^c}$. The week ends at that point, so $a^u = \mathbf{0}$. Otherwise, the binary profile a^u identifies the raisers and the price vector is $\mathbf{p}_j^{a^u}$. Each profile has one entry per chain, equal to one for a cutter or raiser. I denote the resulting scenario-specific vector by $\mathbf{p}_j^{\text{new}}(a^c, a^u)$, or simply $\mathbf{p}_j^{\text{new}}$.

Demand estimates from Section 4 enter the dynamic game as fixed inputs and are not re-estimated. For a transition from $\mathbf{p}_j^{\text{old}}$ to $\mathbf{p}_j^{\text{new}}$, I construct A^+, A^-, and H from the opening-day price change and the new cross-chain price gaps using the seven-day definitions that follow Equation (1). Because no further price change occurs within the week, the opening-day change remains in the seven-day demand history on all seven model days. In the scenario formulas, I suppress the day index on A^+, A^- and H; they are evaluated at the displayed old and new price vectors. Let $g(t) \in \{\text{ordinary}, \text{campaign}, \text{laboratory}\}$ index the three model environments defined above. The demand index satisfies $d(\text{ordinary}) = d(\text{campaign}) = \text{pre}$

and $d(\text{laboratory}) = \text{post}$. Thus, $g(t)$ identifies the model environment and $d(g(t))$ selects its demand coefficients. The fitted mean utility and the daily quantity it implies are

$$\begin{aligned} \widehat{\delta}_{ij}^{g(t)}(\mathbf{p}_j^{\text{new}}, \mathbf{p}_j^{\text{old}}) &= \widehat{\phi}_{ij} - \widehat{\alpha}_{d(g(t))} p_{ij}^{\text{new}} + \widehat{\lambda}_{d(g(t))} L_{ij}(\mathbf{p}_j^{\text{new}}) + \widehat{\gamma}_{\text{inc}} A_{ij}^{+} + \widehat{\gamma}_{\text{cut}} A_{ij}^{-} + \widehat{\gamma}_H H_{ij}, \\ \widetilde{q}_{ij}^{g(t)}(\mathbf{p}_j^{\text{new}}, \mathbf{p}_j^{\text{old}}) &= N_j \frac{\exp\{\widehat{\delta}_{ij}^{g(t)}(\mathbf{p}_j^{\text{new}}, \mathbf{p}_j^{\text{old}})\}}{1 + \sum_{k \in \mathcal{I}} \exp\{\widehat{\delta}_{kj}^{g(t)}(\mathbf{p}_j^{\text{new}}, \mathbf{p}_j^{\text{old}})\}}. \end{aligned} \tag{4}$$

All demand objects follow the definitions in Section 4. Coefficients are reported in Table 3. When evaluating an action, the game sets the common day effect ζ_t and residual demand shock ξ_{ijt} to zero because they capture observed daily variation rather than systematic demand at the scenario prices. If prices do not change, $\mathbf{p}_j^{\text{new}} = \mathbf{p}_j^{\text{old}}$ and $A^+ = A^- = H = 0$.

To compare quantities across drugs, I normalize daily quantity by $\widetilde{N} \equiv \text{median}_j\, N_j$. The drug-specific cost c_j is the median wholesale-cost measure. It is time invariant and common across chains. The μ is complementary basket profit per customer-equivalent unit, measured in thousand CLP.[41] For a given $(a^c, a^u, x_{jt}; \mu)$, normalized daily profit under its action-week price profile is $\overline{\pi}_{ij,t} \equiv [\widetilde{q}_{ij}^{g(t)}(\mathbf{p}_j^{\text{new}}, \mathbf{p}_j^{\text{old}})/\widetilde{N}](p_{ij}^{\text{new}} - c_j + \mu)$.

A complete profile becomes the next price tier.[42] An incomplete profile applies only for the current week; at the opening of the next week, all chains return to the old price tier.

Let $\Delta\overline{\pi}_{ij,t+1}$ denote the change in normalized daily profit next week caused by a price change at the opening of that week. It compares profit at the same next-week prices with and without that price change in the seven-day demand history. The difference can be positive or negative. Within each stationary Bellman block, both terms use that block's fixed demand regime. Both $\overline{\pi}_{ij,t}$ and $\Delta\overline{\pi}_{ij,t+1}$ are evaluated for the $(a^c, a^u, x_{jt}; \mu)$ on the left-hand side below. Chain i's weekly payoff is

$$\pi_{ij}^{g(t)}(a^c, a^u, x_{jt}; \mu) = 7\overline{\pi}_{ij,t} + 7\beta\Delta\overline{\pi}_{ij,t+1}. \tag{5}$$

The current-week term applies for seven days whether a raise is complete or incomplete. A leader's profit nevertheless differs across the two outcomes: only an incomplete profile can leave the leader uniquely high priced, which is captured by H in Equation (4). The observed two-day median response time is a descriptive timing moment and does not shorten the model week.

If any chain cuts, the public price state moves to the destination specified by the appli-

[41]Online Appendix Section B.3, "Monthly losses on below-cost medicine sales," gives the cost construction.

[42]Online Appendix Section B.2, "Construction of price tiers," describes the tier-price construction.

cable cut depth, and all chains charge the resulting lower price profile at the opening of the next week. These next-week price changes determine $\Delta\overline{\pi}_{ij,t+1}$. The term is zero when the price profile does not change at the opening of next week. I include this difference for one week and assume that any price gap is then closed. The event-profit change $\Delta\Pi^{\text{event}}_{ij}(a;\mu)$ in Equation (2), divided by $\widetilde{N}$, is the action-minus-hold difference between the corresponding normalized weekly payoffs in Equation (5), when evaluated at the same price profiles and demand inputs.

5.2 State transitions and public information

States and updating. The price state z_{jt} indexes the persistent three-chain price vector $\mathbf{p}_j^{z_{jt}}$. State I indexes the observed vector in the first model week. A first cut moves the drug to the initial price-war state $0^{(0)}$. The ordinary cut moves one five-percent grid step, $0^{(w)} \to 0^{(w+1)}$; a deeper cut can cross several grid steps. The price grid must contain the destinations implied by the specified cut depths. States 1 and 2 are the constructed Tier 1 and Tier 2 price vectors described in Section 2.[43] No further raise is available from Tier 2.[44]

Here $\mathbf{p}_j^{0^{(w)}}$ is the persistent price vector $\mathbf{p}_j^{z_{jt}}$ when $z_{jt} = 0^{(w)}$. For the initial price-war state, each component of $\mathbf{p}_j^{0^{(0)}}$ is 95 percent of the chain's price in state I. Deeper price-war states satisfy $\mathbf{p}_j^{0^{(w)}} = 0.95^w \mathbf{p}_j^{0^{(0)}}$ for $w = 0, \ldots, K$.

For the ordinary cut depth, the downward-state mapping is $\mathcal{D}(I) = 0^{(0)}$, $\mathcal{D}(0^{(w)}) = 0^{(w+1)}$ for $w = 0, \ldots, K-1$, and $\mathcal{D}(0^{(K)}) = 0^{(K)}$. Thus $\mathcal{D}$ is the ordinary-depth special case of the drug-week mapping $\mathcal{D}_{jt}$ defined below. The upward-state mapping is $\mathcal{U}(z) = 1$ for $z \in \{I, 0^{(0)}, \ldots, 0^{(K)}\}$ and $\mathcal{U}(1) = 2$. Thus $\mathcal{U}$ maps a margin-restoration increase (Tier $0 \to 1$) to state 1 and a rent-extraction increase (Tier $1 \to 2$) to state 2. The post-verification cut transition additionally starts the punishment continuation and is specified separately below.

Price-state transitions. In the Bellman calculation, the next-week grid state follows the within-week sequence in Figure 6. When a drug is reviewed, the chains first choose whether to cut. If at least one chain cuts, the drug moves down the price-war ladder. Write $\mathcal{D}_{jt}$ for the downward mapping at cut depth χ_{jt}. It is the ordinary one-step mapping in Block 1 and incorporates deeper cut destinations in Block 2 and the post-ban games. The same mapping determines action payoffs, Bellman continuation values, and forward price transitions. If all three chains hold, they proceed to the increase stage, where only a complete three-chain

[43] Online Appendix Section B.2, "Construction of price tiers," distinguishes observed benchmark prices from imputed prices and describes the Tier 2 prediction method.

[44] Four drugs attain sustained post-Oficio prices above their Tier 2 benchmarks. Online Appendix Section B.2, "Observed price levels above Tier 2," documents these observations.

increase moves the state up:

$$z_{j,t+1} = \begin{cases} \mathcal{D}_{jt}(z_{jt}), & \text{the drug is reviewed and at least one chain cuts,} \\ \mathcal{U}(z_{jt}), & \text{the drug is reviewed, no chain cuts, and the increase is complete,} \\ z_{jt}, & \text{otherwise.} \end{cases} \tag{6}$$

The cut branch applies except when a drug not already in punishment cuts with $R_t = 1$ at the opening of the week: that deviation starts punishment at $0^{(0)}$ next week; subsequent punishment cuts follow $\mathcal{D}_{jt}$, while verification at the close of a week does not punish that week's cut. The final case includes no review, an all-hold profile, and an incomplete increase. The identity and number of cutters affect current flow payoffs through the cut-price vector $\mathbf{p}_j^{a^c}$, but they do not otherwise affect the transition.

The week-opening indicator R_t records whether the followers know the common procedure. In Game 2, each week with $R_t = 0$ has verification probability η_0: $\Pr(R_{t+1} = 1 \mid R_t = 0) = \eta_0$. Otherwise the followers remain uninformed and continue to play their post-ban-demand version of Game 1 Block 2. Verification is common across drugs, is permanent, and does not itself change any drug's price state. Choices during week t use R_t; a verification at its close changes play from week $t+1$. The leader anticipates this information transition, while uninformed followers do not include the laboratory procedure in their subjective continuation values. Game 1 price changes never change R_t.

Rent extraction participation and belief updating. No outcomes update beliefs in Game 1 or before verification in Game 2. After verification in Game 2, only the outcomes of Tier $1 \to 2$ attempts update beliefs. Each initiated Tier $1 \to 2$ proposal produces one public outcome. Complete participation advances the drug's persistent state from 1 to 2 and counts as a success. An incomplete attempt leaves the state at 1, counts as a failure, and leaves the product eligible for later laboratory reviews. All products selected in week t use the same opening counts, and their outcomes are aggregated when the week closes. Let $n^r_{S,t}$ and $n^r_{F,t}$ denote the numbers of complete and incomplete selected attempts in that week. The count state evolves as

$$(S^r_{t+1}, F^r_{t+1}) = (S^r_t + n^r_{S,t}, F^r_t + n^r_{F,t}). \tag{7}$$

If no Tier $1 \to 2$ product is selected, both weekly counts are zero. Tier $0 \to 1$ outcomes do not enter either count.

At a rent-extraction proposal, the candidate leader's subjective weight on the sequential-following value is

$$m_t^r = m^r(G_t; m_0^r, \nu^r) = \frac{\nu^r m_0^r + S_t^r}{\nu^r + S_t^r + F_t^r}, \tag{8}$$

where m_0^r is the initial subjective weight and ν^r is the prior strength.

Punishment continuation. If any chain cuts after verification, the drug enters the punishment continuation in the following week at the initial price-war state $0^{(0)}$.[45] Raise opportunities are then shut down for that drug. The common cut stage continues to operate. All three chains holding leaves the price state unchanged. A cut applies $\mathcal{D}_{jt}$ and $0^{(K)}$ remains absorbing. The verification state remains one. This drug contributes no further successes or failures to the public counts, but eligible Tier 1 → 2 attempts for other drugs continue to update them. Its value is the separate $V_{ij,t}^{\text{pun}}(z_{jt})$. Public learning does not reopen its increase opportunities.

5.3 Dynamic choices and policy functions

When prices are reviewed. In Game 1, each eligible drug is reviewed with weekly probability ω_1 before the campaign and ω_2 during it. Game 2 uses the review process described in Section 5.1. Before verification, only Tier 0 → 1 increases are available; afterward, both Tier 0 → 1 and Tier 1 → 2 increases are available.[46]

Choices at reached nodes. At a review, the three chains simultaneously choose whether to cut or hold. Before verification, each follower evaluates cutting and holding using Block 2 of Game 1, solved under post-ban demand. It expects the other chains, including Salcobrand, to play that game and does not interpret a price increase as a coordinated proposal. The leader anticipates these follower policies but evaluates cutting and holding using its Game 2 continuation values. After a cut, it can propose a new Tier 0 → 1 increase at a later review. After verification, a cut leads to the punishment continuation in Section 5.2 and ends future increase opportunities for that drug. If all three chains hold in Game 1, they then choose increase or hold simultaneously.

At the Game 2 increase stage, leadership is determined as follows. Before verification, only Salcobrand chooses between leading and waiting on a Tier 0 → 1 increase. It evaluates leading using the followers' Game 1 policies described above and accounts for the probability

[45] Online Appendix A, "(m) Punishment: the threat was a return to the price war," reports the FNE's account of this threat (FNE, 2008, para. 96, pp. 35–36). This account motivates the return to price competition; the timing and transition rules here are model assumptions.

[46] Online Appendix Section D.3, "Forward simulation," describes the review process.

of verification in its next-week continuation value. The followers continue to use those policies and do not interpret the increase as a coordinated proposal. After verification, the model assigns equal probability to each of the six possible orders of the three chains. In that order, each chain chooses between leading and waiting. The first chain that chooses to lead is the unique leader. If all three wait, no increase occurs.

A node history h records the current decision stage and acting chain, the realized chain order when applicable, and all actions already observed in that week's review. At a candidate-leader node it therefore identifies the candidate's position and which earlier candidates have waited; at a follower node it also identifies the leader and any earlier follower action. Simultaneous choices at the current stage are not yet observed. The next equation separates three cases: an increase before verification, a margin-restoration increase after verification, and a rent-extraction increase after verification. Only in the last case does the candidate leader weight the value of sequential follower choices against raising alone. For candidate leader i, the value of leading depends on verification and the proposed tier transition,

$$\bar{v}^{\text{lead}}_{ij,t} = \begin{cases} v^{\text{unverified}}_{ij,t}, & R_t = 0, \quad 0 \to 1, \\ v^{\text{seq}}_{ij,t}, & R_t = 1, \quad 0 \to 1, \\ v^{\text{solo}}_{ij,t} + m^r_t \left(v^{\text{seq}}_{ij,t} - v^{\text{solo}}_{ij,t} \right), & R_t = 1, \quad 1 \to 2. \end{cases} \tag{9}$$

All three values are conditional on reaching the increase stage, after all chains have chosen to hold. Conditional on chain i leading at state x_{jt}, these values are independent of which earlier candidates waited. The sequential branch averages over the two follower orders, so no node-history index is needed. Let a^u denote the three chains' increase profile, with i denoting the candidate leader and k, k' the other two chains. I suppress the known cut profile in the flow payoff. The three values combine weekly profits with discounted continuation values,

$$\begin{aligned} v^{\text{unverified}}_{ij,t} &= \mathbb{E}\Big[\pi^{g(t)}_{ij}(a^u, x_{jt}; \mu) + \beta\big\{(1-\eta_0) V_{ij,t+1}(x^{(0)}_{j,t+1}) + \eta_0 V_{ij,t+1}(x^{(1)}_{j,t+1})\big\} \mid a^u_i = 1, x_{jt}\Big], \\ v^{\text{solo}}_{ij,t} &= \mathbb{E}\big[\pi^{g(t)}_{ij}(a^u, x_{jt}; \mu) + \beta V_{ij,t+1}(x_{j,t+1}) \mid a^u_i = 1,\ a^u_k = 0,\ a^u_{k'} = 0,\ x_{jt}\big], \\ v^{\text{seq}}_{ij,t} &= \mathbb{E}_{\text{seq}}\Big[\pi^{g(t)}_{ij}(a^u, x_{jt}; \mu) + \beta V_{ij,t+1}(x_{j,t+1}) \mid a^u_i = 1, x_{jt}\Big]. \end{aligned} \tag{10}$$

The flow payoff $\pi^{g(t)}_{ij}$ is defined in Equation (5), with $g(t) =$ laboratory and $d(g(t)) =$ post here and the known cut profile omitted. In the unverified value, the expectation is over the other two chains' actions under their post-ban Game 1 Block 2 policies. This value applies only to Salcobrand before verification. The resulting increase profile determines next week's price state as described in Section 5.2. Conditional on that profile, $x^{(0)}_{j,t+1}$ is the next-week state if verification does not occur, with $R_{t+1} = 0$, and $x^{(1)}_{j,t+1}$ is the next-week state if it does, with $R_{t+1} = 1$. These outcomes have probabilities $1 - \eta_0$ and η_0, respectively.

Verification changes next week's information and continuation play, not the current increase profile. The solo value fixes both rivals' actions at hold. The expectation $\mathbb{E}_{\text{seq}}$ uses the followers' sequential Game 2 choices in both specifications, with the first anticipating the second's decision. Each follower uses Equation (12), so participation can be incomplete at either tier. Unit Confidence fixes $m_t^r = 1$ without removing these follower decision nodes. Before verification, the unverified branch uses the uninformed followers' policies in both specifications. Continuation values account for subsequent public-state updates, including learning.

The weight m_t^r affects only the candidate's Tier $1 \to 2$ lead value. it does not change followers' conditional choices. Complete and incomplete Tier $1 \to 2$ attempts update the public counts through Equation (7).

At node h, let $a_i^h \in \{0, 1\}$ denote chain i's choice between taking the relevant action and holding or waiting. The profiles $a^c(h)$ and $a^u(h)$ collect the cut and increase actions resulting from this choice and all remaining within-week choices, denoted a_{rem}^h. These include the other chains' simultaneous choices at the current stage and every later action, including chain i's own later actions. Using the information and beliefs described above, the chain evaluates each action as

$$v_{ij,t}^h(a_i^h; x_{jt}) = \mathbb{E}_{a_{\text{rem}}^h}\left[\pi_{ij}^{g(t)}\big(a^c(h), a^u(h), x_{jt}; \mu\big) + \beta V_{ij,t+1}(x_{j,t+1}) \;\middle|\; a_i^h, x_{jt}, h\right], \tag{11}$$

At a candidate-leader node, the action value in Equation (11) uses the subjective evaluation in Equation (9): $v_{ij,t}^h(1; x_{jt}) = \bar{v}_{ij,t}^{\text{lead}}$. The wait value $v_{ij,t}^h(0; x_{jt})$ integrates over the remaining candidates' lead-or-wait decisions and any ensuing follower decisions, including those of chain i if it is subsequently a follower. After verification, if no candidate remains, waiting leaves all prices unchanged for the week. Before verification, Salcobrand is the only candidate, so its waiting branch contains only the uninformed followers' decentralized choices. Both branches include discounted next-week continuation values. Conditional on reaching a stochastic decision node h, the binary choice follows the rule below. Verified followers use this same sequential logit rule in both specifications.

$$\Pr(a_i^h = 1 \mid x_{jt}, h) = \Lambda\left(\frac{v_{ij,t}^h(1; x_{jt}) - v_{ij,t}^h(0; x_{jt})}{\tau}\right), \tag{12}$$

Λ is the logistic cdf and τ is a fixed action-scale normalization described in Online Appendix Section D.2.

Dynamic value. Let $V_{ij,t}(x_{jt})$ denote chain i's continuation value for drug j outside punishment at the state defined in Equation (3). This value combines current profits with discounted future payoffs. Before verification, the leader evaluates future payoffs using Game 2, whereas the followers expect all three chains to compete as in Block 2 of Game 1, solved under post-ban demand. Both face the same demand, but only the leader understands the procedure. After verification, all three chains evaluate future payoffs using Game 2. The continuation values satisfy

$$V_{ij,t}(x_{jt}) = \mathbb{E}\left[\pi_{ij}^{g(t)}(a^c, a^u, x_{jt}; \mu) + \beta V_{ij,t+1}(x_{j,t+1}) \mid x_{jt}\right] + \mathcal{E}_{ij,t,\tau}(x_{jt}), \tag{13}$$

where the expectation is taken under chain i's information and beliefs about the other chains' behavior. It covers review opportunities, action profiles (a^c, a^u), and the resulting state transitions described in Section 5.2. The model-environment index $g(t)$ and demand index $d(g(t))$ are defined in Section 5.1. Transitions into punishment use $V_{ij,t+1}^{\text{pun}}(z_{j,t+1})$ in place of the ordinary continuation. The term $\mathcal{E}_{ij,t,\tau}$ denotes expected logit surplus at the stochastic decision nodes with action scale τ.[47]

5.4 Estimation, identification, and inference

Both specifications condition on observed laboratory dates and batch capacities.[48]

I calibrate drug-specific campaign cut depths to observed weekly prices and end-of-campaign price levels, while cut choices remain endogenous. Conditional on this calibration, I estimate each specification separately. Unit Confidence estimates four economic parameters, $\boldsymbol{\theta}_U = (\mu, \eta_0, \omega_1, \omega_2)$. These are complementary basket profit μ per customer-equivalent unit, the weekly verification probability while followers remain uninformed η_0, and weekly review probabilities before and during the Game 1 campaign, ω_1 and ω_2. The same μ enters both games within each specification. Unit Confidence fixes the subjective weight at one after verification. Adaptive Confidence jointly estimates $\boldsymbol{\theta}_A = (\mu, \eta_0, \omega_1, \omega_2, m_0^r, \nu^r)$. It adds the initial subjective weight and prior strength for six parameters in total. In the criterion below, $\boldsymbol{\theta}$ denotes the full parameter vector $\boldsymbol{\theta}_A$ or $\boldsymbol{\theta}_U$ for the specification being estimated, with its corresponding search region Θ_F. Follower choices remain endogenous in both specifications.

I set the annual discount factor to 0.80, so $\beta = 0.80^{1/52} \approx 0.9957$ per week.

For each week, I construct two cumulative counts: medicines that have moved beyond Tier 0 and medicines that have reached Tier 2. In simulation, the Tier $0 \to 1$ count records a medicine only when a laboratory-mediated increase in Game 2 succeeds. A complete

[47] This is the expected taste shock to the chosen action under mean-zero Type I extreme-value shocks, weighted by the probability of reaching each stochastic node.

[48] Online Appendix Section D.3, "Forward simulation," describes the construction of review opportunities.

independent increase in Game 1 does not enter this count, even if it advances the medicine's model price state. Each medicine enters each count at most once. In the data, a direct Tier 0 $\rightarrow$ 2 event enters both observed boundary counts but remains one observed price event.

I compare data moments with their simulation means in 9 loss blocks. Within each block, I average the squared standardized discrepancies. The criterion then averages the 9 block losses with equal weight. Let $b = 1, \ldots, 9$ index blocks and $v = 1, \ldots, n_b$ index the statistics within block b, such as weekly prices or price-state shares. For a block containing a single count, $n_b = 1$. Write M_{bv}^{obs} for the observed statistic, $M_{bv,B}^{\text{sim}}(\boldsymbol{\theta}; \tau)$ for its mean across B simulated paths, and σ_{bv} for its fixed scale. Define

$$\mathcal{Q}_{b,B}(\boldsymbol{\theta}; \tau) = \frac{1}{n_b} \sum_{v=1}^{n_b} \left[\frac{M_{bv,B}^{\text{sim}}(\boldsymbol{\theta}; \tau) - M_{bv}^{\text{obs}}}{\sigma_{bv}} \right]^2,$$
$$Q_B(\boldsymbol{\theta}; \tau) = \frac{1}{9} \sum_{b=1}^{9} \mathcal{Q}_{b,B}(\boldsymbol{\theta}; \tau).$$

I minimize this criterion over the bounded parameter search region Θ_F:[49]

$$\widehat{\boldsymbol{\theta}}_B = \arg \min_{\boldsymbol{\theta} \in \Theta_F} Q_B(\boldsymbol{\theta}; \tau).$$

I exclude Game 1 cut-count losses because the observed series contains screened events whereas the simulation records all cut actions. Two cumulative-count blocks compare Tier 0 $\rightarrow$ 1 and Tier 1 $\rightarrow$ 2 medicine counts at five selected weeks. Three price-path blocks compare weekly average prices before the campaign, during the campaign, and in Game 2. The remaining four blocks compare terminal mean price-war depth, the Game 2 boundary price, the Game 2 boundary state distribution, and the verification-time cumulative distribution function (CDF).

For both observed and simulated price-path moments, I first average daily prices within each drug–chain pair and week, so temporary price changes contribute according to their duration. I then take the equally weighted average across drug–chain pairs, in thousand CLP.

Game 2 begins at the start of week 96, carrying forward the price states generated in Game 1. For week 96, I compare observed and simulated weekly average prices. The boundary-state moment compares the simulated end-of-week 96 price-state distribution with the distribution mapped from observed week 96 prices, after the first Game 2 update.

[49]Online Appendix Section D.4, "Estimation and inference," describes the parameter search region.

Under the weekly information transition, verification need not coincide with the first complete laboratory-mediated increase. That observed increase supplies a timing benchmark rather than a direct observation of the followers' information. The verification-time moment compares the share of simulated paths verified by the close of each week with this benchmark indicator, which is zero before week 101 and one thereafter. A transition from $R_t = 0$ to $R_{t+1} = 1$ counts as verification in week t. Informed play starts in week $t + 1$. Thus the week 101 benchmark refers to verification by that week's close, not to informed play at its opening.

I choose the parameters to match the simulated price paths and cumulative numbers of medicines reaching each price tier to their observed counterparts.

The price path and price-state moments discipline μ, ω_1, and ω_2. The parameter μ changes the profit consequence of a price difference. The two review probabilities govern how often a drug can change price in each Game 1 block. Verification timing and early Tier $0 \to 1$ diffusion discipline η_0. The level and evolution of the Tier $1 \to 2$ cumulative-count path also inform the two learning parameters. The initial weight m_0^r affects initiation before rent-extraction outcomes accumulate, while prior strength ν^r governs how strongly subsequent outcomes change that weight in Equation (8). Early participation and its subsequent evolution therefore provide information about the two parameters. I estimate them jointly with the economic parameters using the full criterion.

For Adaptive Confidence, I estimate the parameters conditional on a learning transition prepared before the parameter search. For each candidate parameter vector, I solve the continuation problems and simulate price paths to construct the moments. I use the same moments for both specifications and hold random numbers fixed across parameter candidates.[50]

For inference, I use a laboratory-block bootstrap. Each replication samples laboratories with replacement, retains medicines and their price histories within each sampled laboratory, and re-estimates demand and both dynamic models.

5.5 Results and mechanism assessment

At their separately estimated parameters, the two specifications generate similar Tier $0 \to 1$ paths, while Adaptive Confidence more closely matches the observed cumulative Tier $1 \to 2$ path.

Table 4 reports separately estimated economic parameters for Adaptive Confidence and

[50] Online Appendix Section D.4, "Estimation and inference," describes the numerical search.

Unit Confidence. A common basket-profit parameter enters both games within each specification. Adaptive Confidence jointly estimates the four economic parameters, the initial subjective weight and prior strength. Unit Confidence estimates only the four economic parameters.

Table 4: Conditional dynamic-game parameter estimates

Parameter	Adaptive Confidence	Unit Confidence
Complementary basket profit μ (thousand CLP per customer-equivalent unit)	15.820 (0.539)	16.328 (0.590)
Pre-campaign review probability ω_1 (weekly)	0.02344 (0.00105)	0.02328 (0.00113)
Campaign review probability ω_2 (weekly)	0.10375 (0.00837)	0.10500 (0.00797)
Verification probability η_0 (weekly, given $R_t = 0$)	0.19839 (0.02045)	0.11776 (0.02076)
Initial subjective weight m_0^r (estimated)	0.000725 (0.02011)	–
Prior strength ν^r (effective prior count; estimated)	375 (90.560)	–

Notes: Estimates use 100 simulation paths, the observed laboratory schedule, and fixed calibrated campaign cut depths. Dashes mark learning parameters that do not apply to Unit Confidence. The estimates use approximate game solutions and, for Adaptive Confidence, a learning transition held fixed during estimation. Online Appendix Section D.4, "Estimation and inference," reports the numerical procedure. Parentheses report standard errors from the completed 100-replication paired laboratory-block bootstrap.

Figure 7 compares the two estimated specifications with the observed cumulative tier-transition counts. By week 117, Adaptive Confidence produces 191.58 medicines for Tier 0 $\rightarrow$ 1 and 67.67 for Tier 1 $\rightarrow$ 2, against 208 and 70 in the data. Unit Confidence produces 182.61 and 106.22. Both models begin Tier 0 $\rightarrow$ 1 increases too early. Adaptive Confidence more closely matches the second increase path, whereas Unit Confidence generates more Tier 1 $\rightarrow$ 2 transitions early in Game 2. Because each model has its own estimated economic parameters, the difference between their paths cannot be attributed to the subjective-weight rule alone.

Figure 7: Cumulative price-state transitions

(a) Tier $0 \to 1$ (b) Tier $1 \to 2$

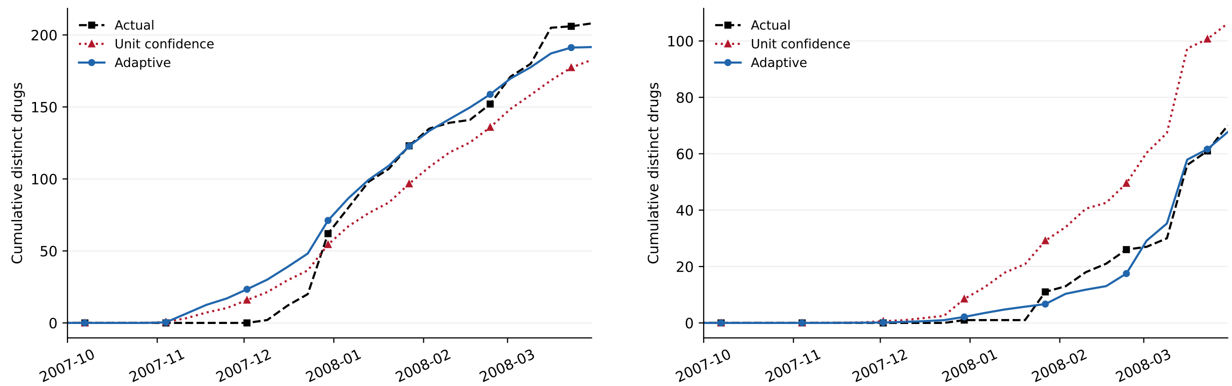

Notes. Each medicine counts once per tier boundary. Lines average 100 simulated paths. Black squares: data; red triangles: Unit Confidence; blue circles: Adaptive Confidence. Simulation starts in January 2007; the displayed window starts in October 2007.

Figure 8 compares weekly cutting and transaction prices. Both models overpredict the week 117 mean price. The predictions are 14.045 thousand CLP under Adaptive Confidence and 13.978 under Unit Confidence, versus 13.732 in the data. Simulated cuts and screened observed events differ in coverage.

Figure 8: Cut timing and average transaction prices

(a) Screened observed events and simulated cut actions (b) Average transaction price

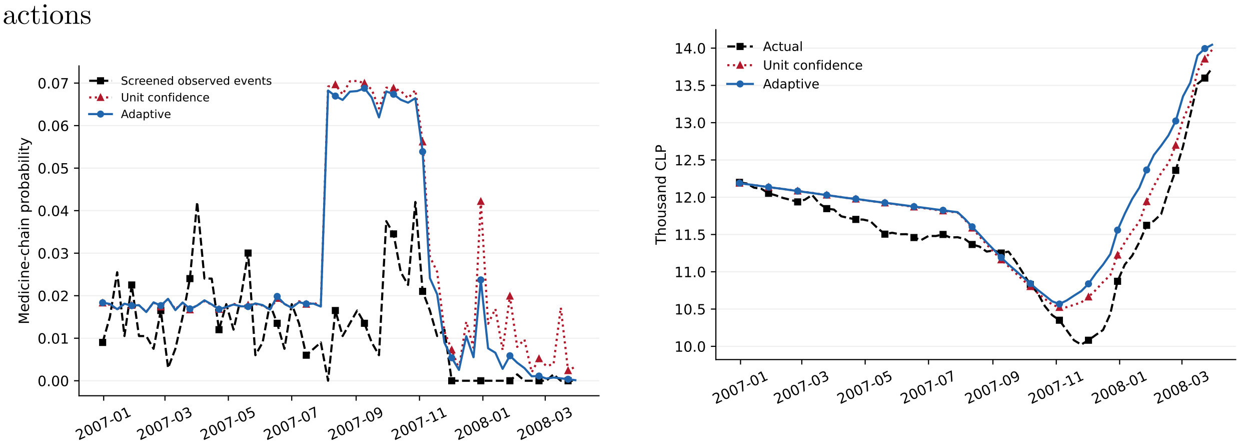

Notes. Panel (a) compares the timing of all simulated cut actions with screened observed cut events; each series is divided by 666 medicine–chain cells, but their levels are not directly comparable because their coverage differs. Panel (b): equally weighted pair-level weekly mean prices, in thousand CLP. Lines average 100 paths; markers follow Figure 7.

Figure 9 varies the subjective weight while holding the economic parameters and approximate learning transition at the Adaptive Confidence estimates in Table 4. It reports the probability that at least one candidate initiates a Tier 1 $\rightarrow$ 2 increase at a verified review with no cut.

Figure 9: Initiation when the subjective weight is low

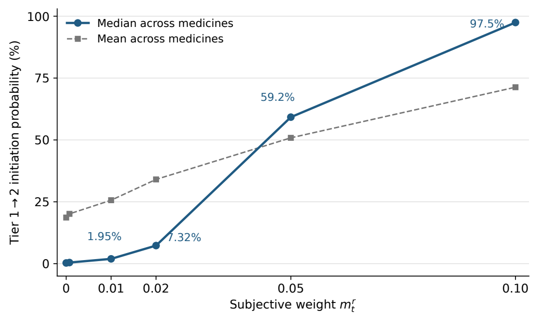

Notes. Mean and median initiation probabilities across 222 equally weighted medicines, conditional on a verified review with no cut. Markers are solved points for $m_t^r \leq 0.10$; connecting lines are visual guides.

Figure 9 shows a sharp increase in initiation probability as the subjective weight rises from 0.01 to 0.05. The median initiation probability is only 1.95 percent at $m_t^r = 0.01$, but rises to 7.32 percent at 0.02 and 59.2 percent at 0.05. In Equation (9), the leader weighs the gain from sequential follower participation relative to raising alone. When that gain is large, even a small weight can make leading attractive. The weight changes the leader's valuation. It is not the actual probability that the followers join.

6 Conclusion

I study how Chile's three largest pharmacy chains moved from a price war to collusion. The demand estimates show that the extra gain from offering a low price fell after the advertisement ban, while the response to ordinary price differences changed little. The ban weakened the incentive to use loss-leader medicines to attract customers, and the sustained price war ended at the same time. The chains then used upstream suppliers to coordinate a sequential price-increase procedure.

I distinguish verification of this procedure from uncertainty about rivals' willingness to participate in further increases. In the model, verification makes the procedure known to the followers. The court record of failed attempts followed by successful repetition motivates this information assumption. Repeated failed attempts preceded the first successful coordinated increase. After that first success, the chains quickly repeated the procedure across medicines to restore positive margins. They subsequently used the same procedure for rent extraction, but these further increases spread more slowly. When considering a rent-extraction increase, the leader still risked losing customers if its rivals did not follow. Knowing how to carry out the procedure did not settle whether the followers would find another increase profitable.

I compare Adaptive Confidence, which updates the weight the leader places on the expected profit from sequential follower responses relative to raising alone, with Unit Confidence, which fixes that weight at one. Complete rent-extraction attempts raise the weight and incomplete attempts lower it. Followers choose whether to join according to their incentives in both models. Adaptive Confidence is closer to the observed spread of rent extraction. By the end of the estimation period, it puts 68 medicines at the rent-extraction tier against 70 observed. Knowing how to coordinate a price increase does not establish that rivals will join it.

References

17º Juzgado Civil de Santiago. Medida precautoria: Farmacias Ahumada S.A. con Farmacias Cruz Verde S.A. — Ley 20,169 de competencia desleal. 17º Juzgado Civil de Santiago, 6 November 2007. Reported by CIPER Chile, "El dossier del caso farmacias" (9 April 2009)., 2007.

Victor Aguirregabiria and Jihye Jeon. Firms' beliefs and learning: Models, identification, and empirical evidence. *Review of Industrial Organization*, 56(2):203–235, 2020.

Victor Aguirregabiria and Arvind Magesan. Identification and estimation of dynamic games when players' beliefs are not in equilibrium. *The Review of Economic Studies*, 87(2): 582–625, 2020.

Jorge Alé-Chilet. Gradually Rebuilding a Relationship: The Emergence of Collusion in Retail Pharmacies in Chile. Working paper, Bar-Ilan University, 2016.

Jorge Alé-Chilet. Collusive Price Leadership in Retail Pharmacies in Chile. Working paper, Bar-Ilan University, 2018.

Jorge Alé-Chilet and Juan Pablo Atal. Trade associations and collusion among many agents: Evidence from physicians. *RAND Journal of Economics*, 51(4):1197–1221, 2020.

Joe S. Bain. *Industrial Organization*. John Wiley and Sons, New York, 1959.

Lee Benham. The effect of advertising on the price of eyeglasses. *Journal of Law and Economics*, 15(2):337–352, 1972.

David P Byrne and Nicolas De Roos. Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109(2):591–619, 2019.

David P. Byrne, Nicolas de Roos, Matthew S. Lewis, Leslie M. Marx, and Xiaosong Wu. Negotiating coordination: A study in retail gasoline. Technical report, SSRN, August 2026. URL https://ssrn.com/abstract=6338858. Last revised August 27, 2026.

John F. Cady. An estimate of the price effects of restrictions on drug price advertising. *Economic Inquiry*, 14(4):493–510, 1976.

Carla Castillo-Laborde and Pablo Villalobos Dintrans. Caracterización del gasto de bolsillo en salud en Chile: una mirada a dos sistemas de protección. *Revista Médica de Chile*, 141(11):1456–1463, 2013. doi: 10.4067/S0034-98872013001100013. URL https://www.scielo.cl/scielo.php?pid=S0034-98872013001100013&script=sci_arttext.

Central de Abastecimiento del Sistema Nacional de Servicios de Salud. Balance de gestión integral, año 2007. Technical report, Ministerio de Salud, Chile, 2008. URL https://www.dipres.gob.cl/597/articles-36949_doc_pdf.pdf.

Daniel Chaves and Marco Duarte. The inner workings of a hub-and-spoke cartel in the automotive fuel industry. *American Economic Journal: Microeconomics*, 17(1):41–65, 2025.

Zhijun Chen and Patrick Rey. Loss leading as an exploitative practice. *American Economic Review*, 102(7):3462–3482, 2012.

David Debrott. Gasto de bolsillo en salud en Chile 2005: Estructura, participación relativa y tendencias. *Revista Chilena de Salud Pública*, 11(1):38–42, 2007. URL https://revistasaludpublica.uchile.cl/index.php/RCSP/article/view/38913.

Patrick DeGraba. Volume discounts, loss leaders and competition for more profitable customers. Working Paper 260, Federal Trade Commission, Bureau of Economics, 2003.

Ulrich Doraszelski, Gregory Lewis, and Ariel Pakes. Just starting out: Learning and equilibrium in a new market. *American Economic Review*, 108(3):565–615, 2018.

Gonzalo Durán and Marco Kremerman. Informe retail: Capítulo farmacias. Technical report, Fundación SOL, December 2007. URL https://fundacionsol.cl/cl_luzit_herramientas/static/wp-content/uploads/2010/09/Cuaderno-3-Farmacias.pdf. Prepared for the Departamento de Estudios, Dirección del Trabajo de Chile.

FNE. Requerimiento contra farmacias ahumada s.a., cruz verde s.a. y salcobrand s.a. Requerimiento rol c no. 184-08, Fiscalía Nacional Económica, Chile, 2008. URL https://www.fne.gob.cl/wp-content/uploads/2011/03/requ_0009_2008.pdf. Legal complaint, Tribunal de Defensa de la Libre Competencia.

David Genesove and Wallace P. Mullin. The sugar institute learns to organize information exchange. In Naomi R. Lamoreaux, Daniel M. G. Raff, and Peter Temin, editors, *Learning by Doing in Markets, Firms, and Countries*, pages 103–144. University of Chicago Press, Chicago, 1999.

David Genesove and Wallace P Mullin. Rules, communication, and collusion: Narrative evidence from the sugar institute case. *American Economic Review*, 91(3):379–398, 2001.

Heiko A. Gerlach and Lan Nguyen. Price staggering in cartels. *International Journal of Industrial Organization*, 77:102757, 2021. doi: 10.1016/j.ijindorg.2021.102757.

Ronald L. Goettler and Brett R. Gordon. Does AMD spur Intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200, 2011. doi: 10.1086/664615.

Edward J. Green and Robert H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100, 1984.

Edward J Green, Robert C Marshall, and Leslie M Marx. Tacit collusion in oligopoly. *The Oxford handbook of international antitrust economics*, 2:464–497, 2014.

Joseph E. Harrington. A theory of collusion with partial mutual understanding. *Research in Economics*, 71(1):140–158, 2017. doi: 10.1016/j.rie.2016.11.005.

Jr. Harrington, Joseph E. Lectures on Collusive Practices. CRESSE lecture, University of Pennsylvania, 2018.

Yufeng Huang, Paul B Ellickson, and Mitchell J Lovett. Learning to set prices. *Journal of Marketing Research*, 59(2):411–434, 2022.

Mitsuru Igami and Takuo Sugaya. Measuring the incentive to collude: the vitamin cartels, 1990–99. *The Review of Economic Studies*, 89(3):1460–1494, 2022.

Jihye Jeon. Learning and investment under demand uncertainty in container shipping. *The RAND Journal of Economics*, 53(1):226–259, 2022.

Rajiv Lal and Carmen Matutes. Retail pricing and advertising strategies. *Journal of Business*, 67(3):345–370, 1994.

Nathan H Miller, Gloria Sheu, and Matthew C Weinberg. Oligopolistic price leadership and mergers: The united states beer industry. *American Economic Review*, 111(10):3123–59, 2021.

Jeffrey Milyo and Joel Waldfogel. The effect of price advertising on prices: Evidence in the wake of 44 Liquormart. *American Economic Review*, 89(5):1081–1096, 1999.

Servicio Nacional del Consumidor. Qué hay detrás de las promesas publicitarias. SERNAC market report, 6 June. URL: https://www.sernac.cl/portal/619/w3-article-6130.html, 2006.

TDLC. Sentencia no. 119/2012: Farmacias ahumada, cruz verde y salcobrand. Sentencia 119/2012, Tribunal de Defensa de la Libre Competencia, Chile, 2012. URL https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia_119_2012.pdf.

Øyvind Thomassen, Howard Smith, Stephan Seiler, and Pasquale Schiraldi. Multi-category competition and market power: A model of supermarket pricing. *American Economic Review*, 107(8):2308–2351, 2017.

Chenyu Yang. Vertical structure and innovation: A study of the SoC and smartphone industries. *The RAND Journal of Economics*, 51(3):739–785, 2020. doi: 10.1111/1756-2171.12339.